ALCATEL.LUCENT Enterprise Business Group

IP Networking Portfolio

Network Solutions with

OmniSwitch 9000 Series

Technical Document

Table of Contents

OmniSwitch 9000 Series

Introduction

As Enterprises search for competitive advantages in the market place and become increasingly dependent on their networks to conduct business, new network requirements have rapidly emerged, exceeding the capabilities of successive technological advancements.

Enterprise new challenges include:

- Highly available, highly secure, highly intelligent, highly manageable and highly scalable Enterprise networks
- The rapid growth of Internet, Intranet and Extranet networking requirements
- Emerging new applications: converged IP applications, streaming media, desktop conferencing, IP-storage, etc.
- Increased clients (vendors, partners, customers, distributors, telecommuters, etc.) access to network resources
- Support high-density traffic aggregation in mission critical business network cores
- Today's Enterprise networks demanding higher switching capacities to improve performance and to accommodate higher 10GigE port densities. The trends in this market are mostly price driven.
  - 10GigE – Performance requirements
  - 10GigE – Port density requirements
- Various government requirements for IPv6
- Requirements for fast network response times

To meet these new market demands, the solution is to provide intelligent devices capable of supporting a host of advanced features for high volume intelligent traffic handling and processing. **Intelligent performance is essential.**

OmniSwitch 9000 Series value propositions, to support the Enterprise new challenges include:

- **Value**
- **High Availability**
- **Embedded Security**
- **Distributed Intelligence**
- **Simplified Manageability**

The OmniSwitch 9000 Family

The OmniSwitch 9000 platforms provide high availability, embedded security, distributed intelligence, easy-to-manage, high performance, and high throughput designed mainly for enterprise core networks.
These features are available in a compact form factor at an extremely aggressive price point 4.2.1
The award winning Alcatel OmniSwitch 9000 family (OS9000s) offers a range of full-featured, high-performance modular-based configuration, triple-speed (10/100/1000Mbps) Ethernet, Gigabit Ethernet, and 10-Giagbit Ethernet switches, including a low-cost entry point chassis. The OS9000s deliver future-proof solutions with advanced security and QoS features for use in small-to-large enterprise cores, in the aggregation layer and in wiring closets with flexible power-over-Ethernet support. The OS9000s are a part of Alcatel's end-to-end enterprise switch family.
The Alcatel OmniSwitch 9000s are designed to anticipate future network needs with wire-rate processing for IPv4/IPv6 and support for unicast and multicast applications such as voice-over-IP and video collaboration. The switches support future edge requirements as Gigabit Ethernet to the desktop becomes commonplace and demand for power-over-Ethernet (PoE) capability increases.
They provide wire rate layer-2 forwarding and layer-3 routing with advanced intelligent services. The OS9000 switches increase network performance, improve application response times, secure the LAN and enhance user productivity by maximizing network capacity and services over existing category 5/5E/6 cabling. With triple-speed (10/100/1000Mbps) Ethernet NI Modules, Gigabit Ethernet NI Modules, 10-Giagbit Ethernet capabilities NI Modules, and IEEE 802.3af PoE support on triple-speed (10/100/1000Mbps) Ethernet Modules, the small to large enterprises can now simultaneously protect their current investment in legacy end devices while providing for seamless migration in the future.

- A resilient, affordable & high performance solution
- Large Gigabit & 10 Gigabit Ethernet port density
  The 18-slot OS9800 offers up to 768 Copper Gigabit Ethernet ports, 384 Optical Fiber Gigabit Ethernet ports and up to 96 10-Gigabit Ethernet ports.
  The 10-slot OS9700 offers up to 384 Copper Gigabit Ethernet ports, 192 Optical Fiber Gigabit Ethernet ports and up to 48 10-Gigabit Ethernet ports. 4.2.10
  The 5-slot OS9600 offers up to 192 Copper Gigabit Ethernet ports, 96 Optical Fiber Gigabit Ethernet ports and up to 24 10-Gigabit Ethernet ports.
- Redundant architecture for converged networks
- Native support for IPv4 & IPv6 for network future proofing
- A totally new Architecture
- Extensive Multicast support (L2/IPv4/IPv6)
- Enhanced network response time
- Protecting the control plane from external attacks (DoS)

Alcatel.Lucent AOS Releases 6.1.1r01 & 6.1.1r02 & 6.1.3r01 supports the following hardware:

- OmniSwitch 9600 & 9700 & 9800 Chassis Models
  - OS9600: 5 slots chassis (4 NIs + 1 CMM)
  - OS9700: 10 slots chassis (8 NIs + 2 CMMs)
  - OS9800: 18 slots chassis (16 NIs + 2 CMMs
- OS9600-CMM & OS9700-CMM & OS9800-CMM Chassis Management Modules
- OS9-GNI-C24 (triple-speed Ethernet copper Module)
- OS9-GNI-P24 (triple-speed Ethernet copper Module with PoE)
- OS9-GNI-U24 (Gigabit Ethernet Fiber Module)
- OS9-XNI-U2 (2 x 10-Gigabit Ethernet Fiber Module)
- OS9-XNI-U6 (6 x 10-Gigabit Ethernet Fiber Module)

The OmniSwitch 9000 Series product family with a complete and state-of-the-art set of industry-based features is perfect for the following applications:

- Enterprise Edge deployments & branch offices
- L3 Aggregation / distribution layer switches in three-tier networks
- Medium to large enterprise core switching
- Quality of service (QoS) for mission critical applications
- Data center server clusters

OmniSwitch 9600

The OmniSwitch 9600 (OS9600) is Alcatel's low-cost, entry-point LAN switch solution that provides the small enterprise the best performance-to-price available today. It also future proofs the investment since the network can be expanded to support a larger core inexpensively by reusing its current OS9000 cards in the larger chassis offered by the OS9700 and OS9800. By offering the same features and capabilities as the other OmniSwitch 9000 switches, more enterprise networks are able to afford a core that supports converged voice, video and data and other applications.
The Alcatel OmniSwitch 9600 is a five-slot chassis supporting one chassis management module (CMM) and four network interface (NI) modules. It offers a wide range of GigE and 10GigE interfaces providing the industry's most flexible combination of Ethernet interfaces for use in a wiring closet. It also offers power-over-Ethernet to support IP telephones, WLAN access points and video cameras. VoIP and video performance is also enhanced in an OmniSwitch-based network through the use of policy-based QoS using OmniVista NMS PolicyView.

OmniSwitch 9700

The OmniSwitch 9700 is a high-density chassis with two slots for control and 8 slots for network interfaces supporting an aggregation of up to 384 Copper GigE ports, 192 Optical Fiber GigE ports or 48 10GigE ports. Designed for smart continuous switching operation, the two center slots are dedicated to chassis management modules (CMMs) allowing redundant configurations. CMMs provide two critical functions – active/standby resiliency for control and active/active redundancy for the switching fabric. The Alcatel OmniSwitch 9700 was recently named as Network Computing Editor's Choice for its superb manageability and ease of use.

OmniSwitch 9800

For those applications where a much larger port density is required, the OS9800 doubles the OS9700 capabilities with 16 available slots supporting an aggregation of up to 768 Copper GigE ports, 384 Optical Fiber GigE ports or 96 10GigE ports along with two slots for control (Plus, the OS9800 has subcomponents such as the power supply unit (PSU), fan tray and network interface cards (NICs) that are all compatible and interchangeable with other OS9000s, reducing the cost of keeping spares and lowering total cost of ownership.)

Hardware Overview

The OmniSwitch 9000 switches are available in three chassis configurations—the 18-slot OmniSwitch 9800 (OS9800), the 10-slot OmniSwitch 9700 (OS9700), and the 5-slot OmniSwitch 9600 (OS9600).
The 18-slot OS9800 offers up to 384 Gigabit Ethernet ports and up to ninety-six 10-Gigabit Ethernet ports. The 10-slot OS9700 offers up to 192 Gigabit Ethernet ports and up to fourty-eight10-gigabit Ethernet ports while the 5-slot OS9600 offers up to 96 Gigabit Ethernet ports and up to twenty-four 10-gigabit Ethernet ports.

The OmniSwitch 9800

The OmniSwitch 9800 is a high performance switch offering 16 slots for Gigabit Ethernet and/or 10-Gigabit Ethernet Network Interface (NI) modules. An additional two slots are reserved for primary and redundant Chassis Management Modules (CMMs). The OmniSwitch 9800 supports a maximum of four power supplies.
Note. Power supply requirements are based on the number of NIs installed in the chassis.

The following illustrations outline the major components of the OmniSwitch 9800 switch.

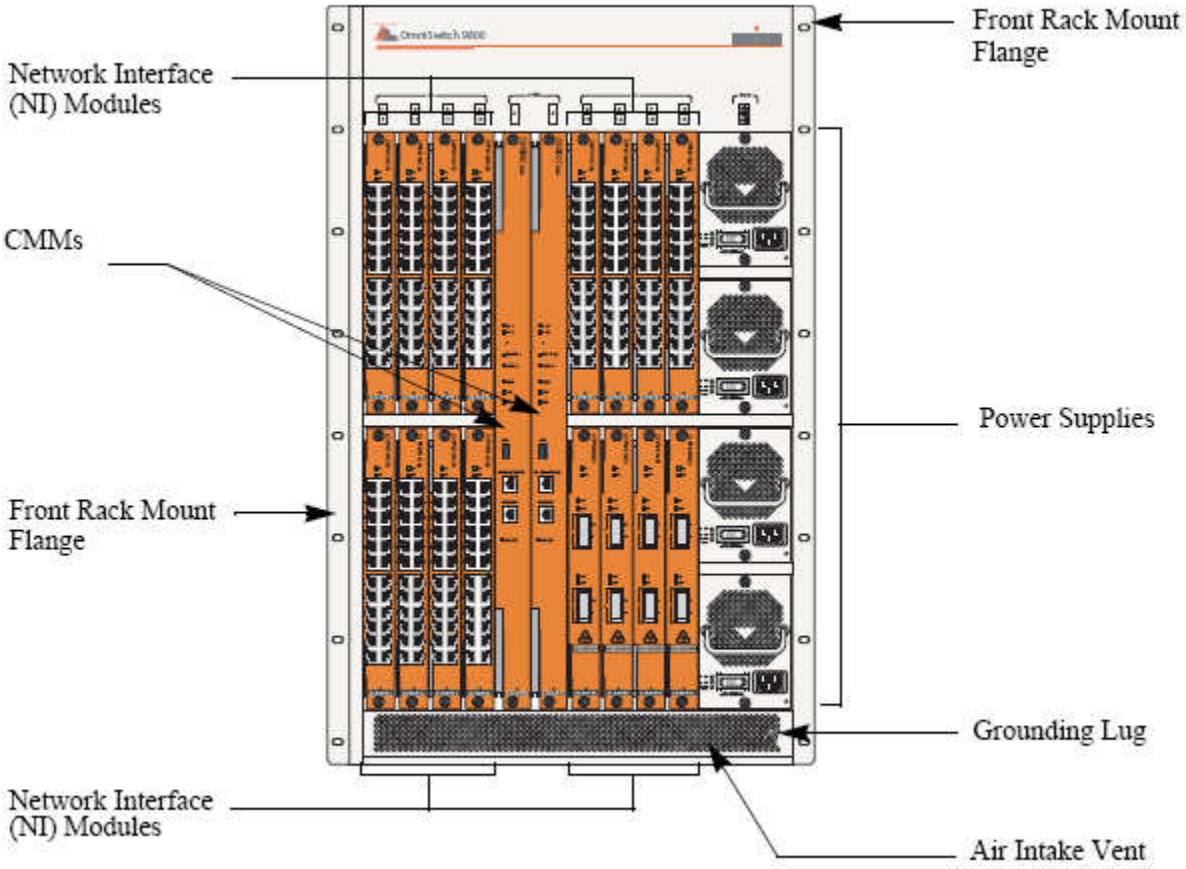

OmniSwitch 9800 Front View

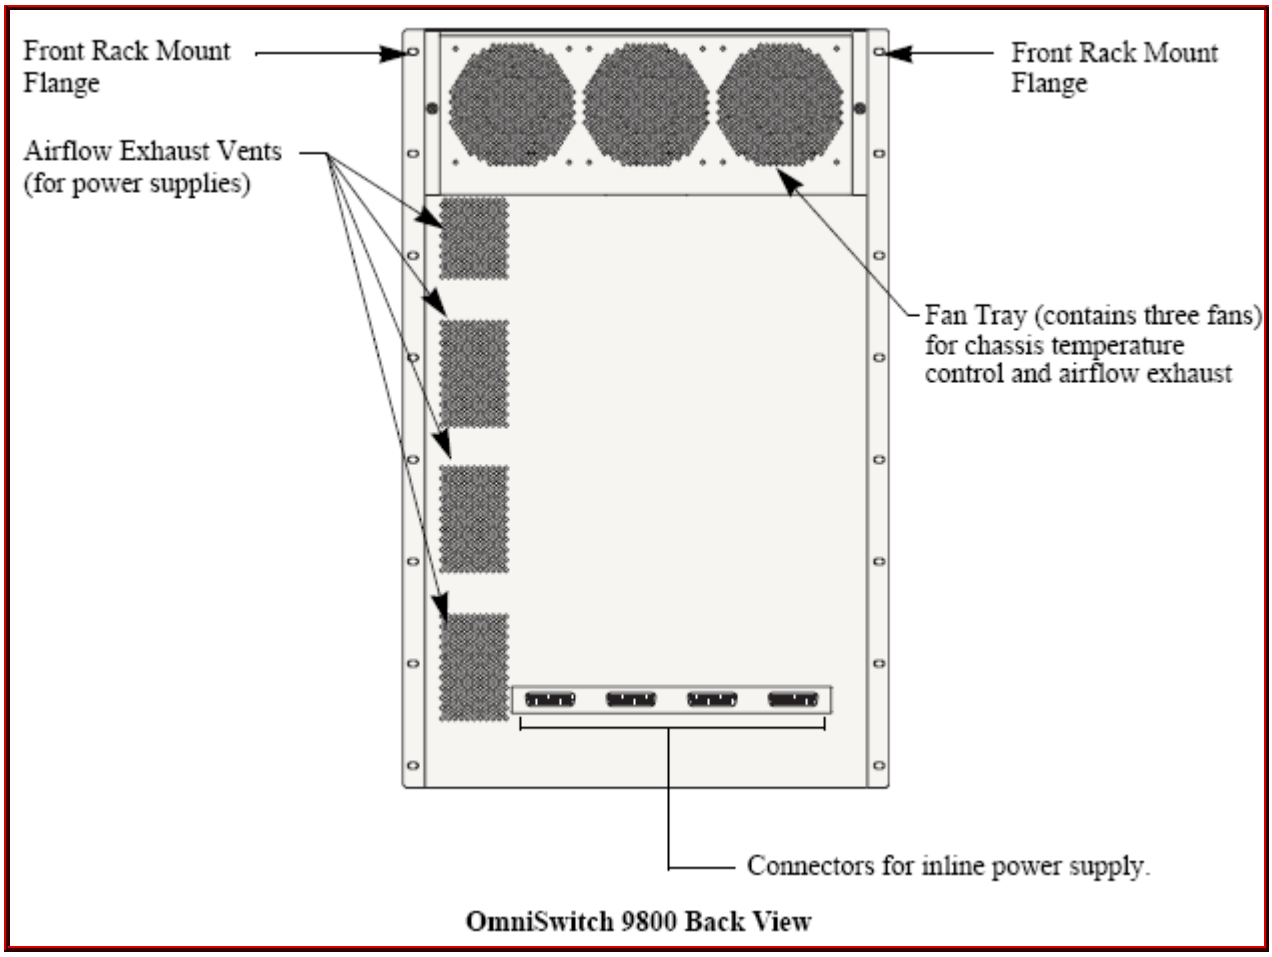

OmniSwitch 9800 Back View

OmniSwitch 9800 Chassis Technical Specifications	
OmniSwitch 9800	The OmniSwitch 9800 is a 18-slot large enterprise core switch. The OmniSwitch 9800 offers up to 384 Gigabit Ethernet ports and can also be equipped with up to 32/96 (Fully-populated with 16 x OS9-XNI-U2 and/or OS9-XNI-U6 modules, with each XNI containing two or six XFPs) 10-Gigabit Ethernet ports. The OmniSwitch 9800 chassis contains the following major components: • Sixteen Network Interface (NI) module slots • Two Chassis Management Module (CMM) slots • Power supply bay holding up to four power supplies • Fan tray with three fans
Total slots per chassis	18
Total slots for network interface (NI) modules	16
Total slots for Chassis Management Module (CMM)	2
Total bays for power supplies	4
Total 10/100/1000BASE-T copper 10/100/1000Mbps Ethernet ports available	384 (Fully populated with OS9-GNI-C24). No other NI module types installed.
Total 10/100/1000BASE-T copper 10/100/1000Mbps Ethernet ports with PoE available	384 (Fully-populated with 16 x OS9-GNI-P24 modules) No other NI module types installed.
Total 1000BASE-X fiber Gigabit Ethernet ports available	384 (Fully populated with OS9-GNI-U24 modules). No other NI module types installed.
Total 10GBASE-X 10-Gigabit Ethernet ports available	**Wire-rate:** 32 (Fully populated with OS9-XNI-U2 modules, with each XNI containing two XFPs. No other NI modules installed.) **Oversubscribed:** 96 (Fully populated with OS9-XNI-U6 modules, with each XNI containing six XFPs. No other NI modules installed.) **The oversubscription ratio is: 2.5:1**
Power Consumption	OS9800-Chassis & Fans = 80W
Full Duplex support	Full duplex is supported on Gigabit Ethernet ports and 10-Gigabit Ethernet ports.
Environmental Requirements	OmniSwitch 9000 Series switches have the following environmental and airflow requirements: • The installation site must maintain a temperature between 0° and 45° Celsius (32° and 113° Fahrenheit) and not exceed 95 percent maximum humidity (non-condensing) at any time. • Be sure to allow adequate room for proper air ventilation at the front, back, and sides of the switch. No clearance is necessary at the top or bottom of the chassis.
Electrical Requirements	OmniSwitch 9000 Series switches have the following general electrical requirements: • Each switch requires one grounded electrical outlet for each power supply installed in the chassis (up to four for OS9800 switches, up to three for OS9700 switches; up to two for OS9600 switches). OmniSwitch 9000 Series switches offer both AC and DC power supply support. For switches using AC power connections, each supplied AC power cord is 2 meters (approximately 6.5 feet) long. Do not use extension cords. **Redundant AC Power.** It is recommended that each AC outlet reside on a separate circuit. With redundant AC, if a single circuit fails, the switch's remaining power supplies (on separate circuits) are likely to remain unaffected and can, therefore, continue operating.
OmniSwitch 9800 Chassis Dimensions	
Overall Width (including rack-mount flanges)	19 1/8 inches
Chassis Width (rack-mount flanges not included)	17 9/16 inches
Height	29 1/4 inches
Height (rack units)	17 RU
Overall Depth (including required fan tray)	17 5/16 inches
Chassis Depth (fan tray not included)	14 3/4 inches

OmniSwitch 9700

The OmniSwitch 9700 is a high performance switch offering eight slots for Ethernet, Gigabit Ethernet, and/or 10Gigabit Ethernet Network Interface (NI) modules. Additional two slots are reserved for primary and redundant Chassis Management Modules (CMMs). The OmniSwitch 9700 supports a maximum of three power supplies.
Note. Power supply requirements are based on the number of NIs installed in the chassis.

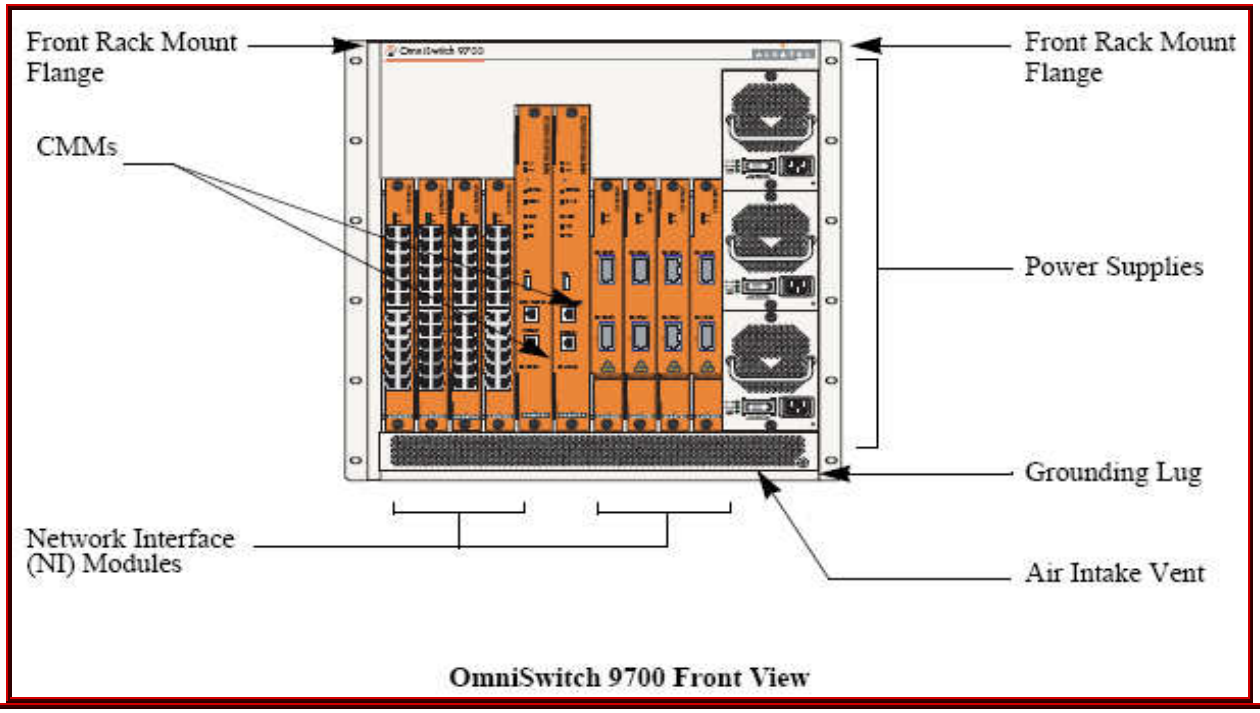

OmniSwitch 9700 Front View

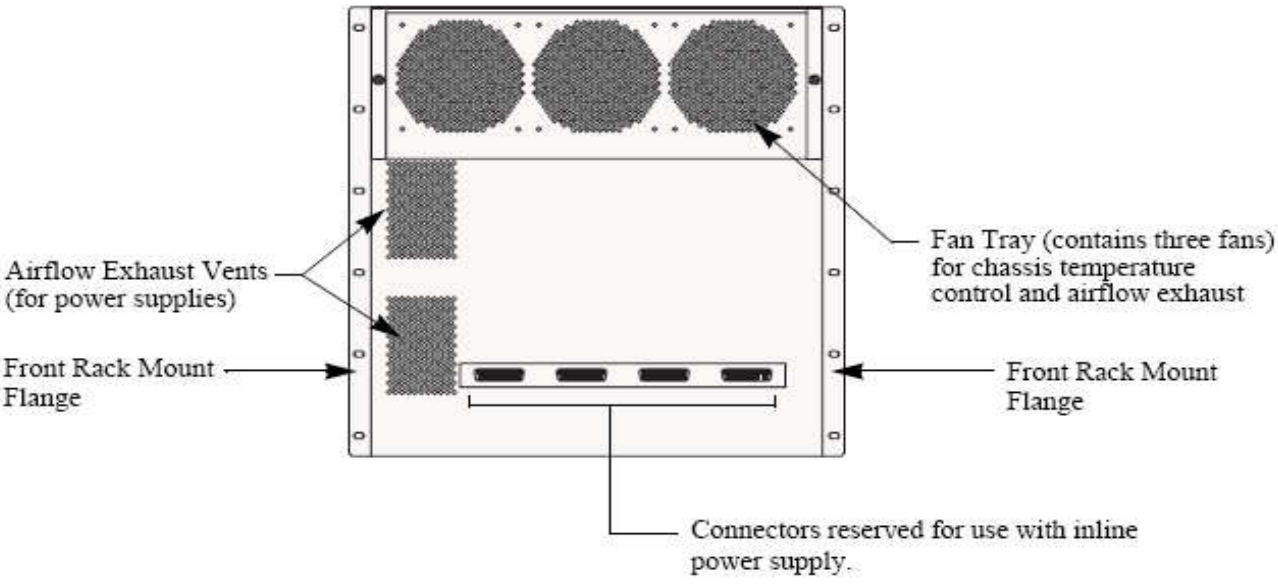

OmniSwitch 9700 Back View

OmniSwitch 9700 Chassis Technical Specifications	
OmniSwitch 9700	The OmniSwitch 9700 is a 10-slot large enterprise core switch. The OmniSwitch 9700 offers up to 192 Gigabit Ethernet ports and can also be equipped with up to 16/48 (Fully-populated 8 x OS9-XNI-U2/OS9-XNI-U6 modules, with each XNI containing two or six XFPs) 10-Gigabit Ethernet ports. The OmniSwitch 9700 chassis contains the following major components: • Eight Network Interface (NI) module slots • Two Chassis Management Module (CMM) slots • Power supply bay holding up to three power supplies • Fan tray with three fans
Total slots per chassis	10
Total slots for network interface (NI) modules	8
Total slots for Chassis Management Module (CMM)	2
Total bays for power supplies	3
Total 10/100/1000BASE-T copper 10/100/1000Mbps Ethernet ports available	192 (Fully-populated with 8 x OS9-GNI-C24 modules) No other NI module types installed.
Total 10/100/1000BASE-T copper 10/100/1000Mbps Ethernet ports with PoE available	192 (Fully-populated with 8 x OS9-GNI-P24 modules) No other NI module types installed.
Total 1000BASE-X fiber Gigabit Ethernet ports available	192 (Fully-populated with 8 x OS9-GNI-U24 modules) No other NI module types installed.
Total 10GBASE-X 10-Gigabit Ethernet ports available	**Wire-rate:** 16 (Fully-populated with 8 OS9-XNI-U2 modules, with each XNI containing two XFPs.) **Oversubscribed:** 48 (Fully-populated with 8 OS9-XNI-U6 modules, with each XNI containing six XFPs.) **The oversubscription ratio is: 2.5:1**
Power Consumption	OS9700-Chassis & Fans = 80W
Full Duplex support	Full duplex is supported on Gigabit Ethernet ports and 10-Gigabit Ethernet ports.
Environmental Requirements	OmniSwitch 9000 Series switches have the following environmental and airflow requirements: • The installation site must maintain a temperature between 0° and 45° Celsius (32° and 113° Fahrenheit) and not exceed 95 percent maximum humidity (non-condensing) at any time. • Be sure to allow adequate room for proper air ventilation at the front, back, and sides of the switch. No clearance is necessary at the top or bottom of the chassis.
Electrical Requirements	OmniSwitch 9000 Series switches have the following general electrical requirements: • Each switch requires one grounded electrical outlet for each power supply installed in the chassis (up to three for OS9700 switches; up to two for OS9600 switches). OmniSwitch 9000 Series switches offer both AC and DC power supply support. For switches using AC power connections, each supplied AC power cord is 2 meters (approximately 6.5 feet) long. Do not use extension cords. **Redundant AC Power.** It is recommended that each AC outlet reside on a separate circuit. With redundant AC, if a single circuit fails, the switch's remaining power supplies (on separate circuits) are likely to remain unaffected and can, therefore, continue operating.
OmniSwitch 9700 Chassis Dimensions	
Overall Width (including rack-mount flanges)	19 1/8 inches
Chassis Width (rack-mount flanges not included)	17 9/16 inches
Height	19 1/4 inches
Height (rack units)	11 RU
Overall Depth (including required fan tray)	17 5/16 inches
Chassis Depth (fan tray not included)	14 3/4 inches

OmniSwitch 9600

The OmniSwitch 9600 is a high performance switch offering four slots for Gigabit Ethernet and/or 10- gigabit Ethernet Network Interface (NI) modules. An additional one slot is reserved for the primary Chassis Management Module (CMM). The OmniSwitch 9600 supports a maximum of two load sharing power supplies on the front panel and there are optional power entry provisions, which consist of three DB-25 connectors mounted on the rear panel of the chassis for PoE applications. We can use either OS9-IPSHELF or 360W/510W power supplies. The first two connectors support OS9-IP-SHELF power supplies and the third connector support 360W/510W (aka. 230W/390W) power supplies.
Note. Power supply requirements are based on the number of NIs installed in the chassis.

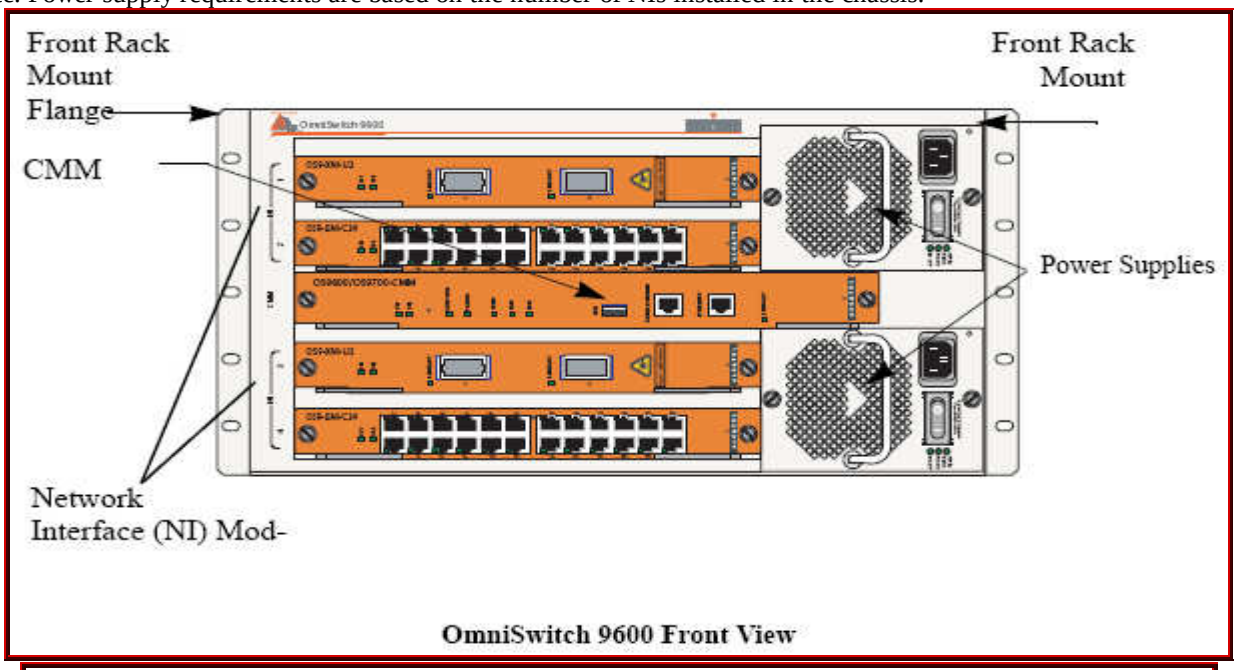

OmniSwitch 9600 Front View

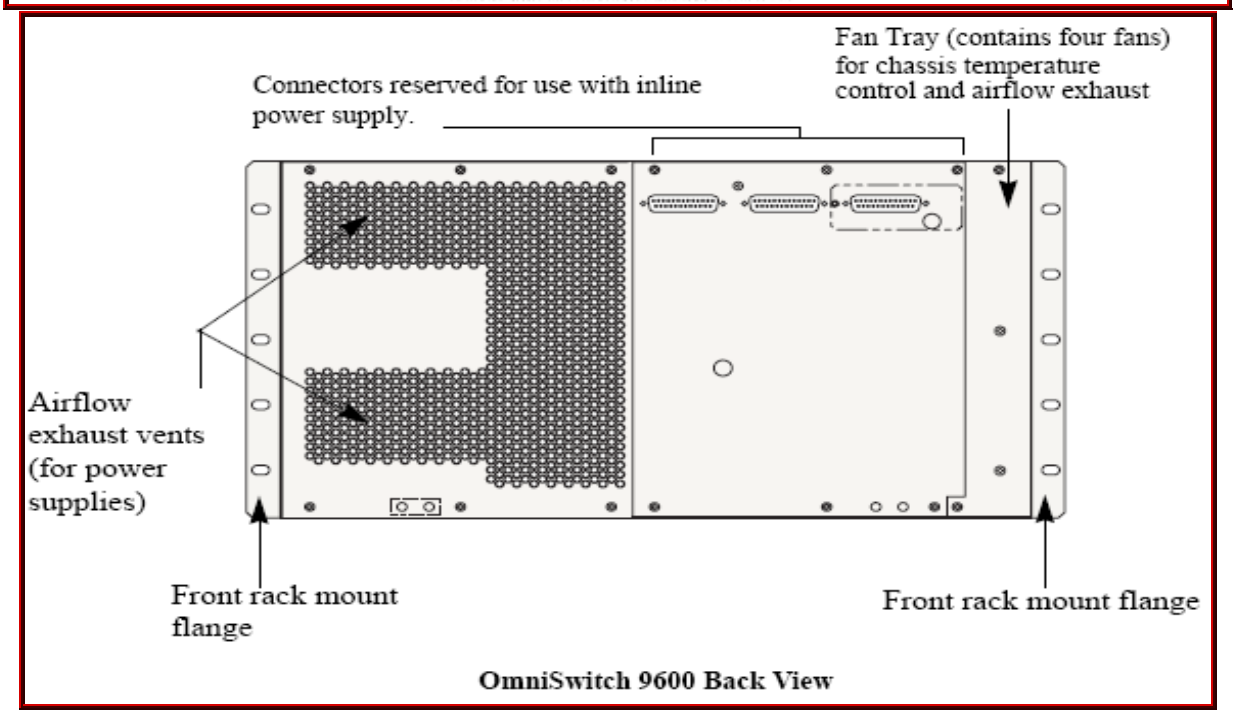

OmniSwitch 9600 Back View

OmniSwitch 9600 Chassis Technical Specifications	
OmniSwitch 9600	The OmniSwitch 9600 is a 5-slot large enterprise core switch. The OmniSwitch 9600 offers up to 96 Gigabit Ethernet ports and can also be equipped with up to 8/24 (Fully-populated with 4 x OS9-XNI-U2/OS9-XNI-U6 modules, with each XNI containing two/six XFPs) 10-Gigabit Ethernet ports. The OmniSwitch 9600 chassis contains the following major components: • Four Network Interface (NI) module slots • One Chassis Management Module (CMM) slots • Power supply bay holding up to two power supplies • Fan tray with four fans
Total slots per chassis	5
Total slots for network interface (NI) modules	4
Total slots for Chassis Management Module (CMM)	1
Total bays for power supplies	2
Total 10/100/1000BASE-T copper 10/100/1000Mbps Ethernet ports available	96 (Fully-populated with 4 x OS9-GNI-C24 modules)
Total 10/100/1000BASE-T copper 10/100/1000Mbps Ethernet ports with PoE available	96 (Fully-populated with 4 x OS9-GNI-P24 modules)
Total 1000BASE-X fiber Gigabit Ethernet ports available	96 (Fully-populated with 4 x OS9-GNI-U24 modules)
Total 10GBASE-X 10-Gigabit Ethernet ports available	**Wire-rate:** 8 (Fully-populated with 4 OS9-XNI-U2 modules, with each XNI containing two XFPs.) **Oversubscribed:** 24 (Fully-populated with 4 OS9-XNI-U6 modules, with each XNI containing six XFPs.) **The oversubscription ratio is: 2.5:1**
Power Consumption	OS9600-Chassis & Fans = 42watts
Full Duplex support	Full duplex is supported on Gigabit Ethernet ports and 10-Gigabit Ethernet ports.
Environmental Requirements	OmniSwitch 9000 Series switches have the following environmental and airflow requirements: • The installation site must maintain a temperature between 0° and 45° Celsius (32° and 113° Fahrenheit) and not exceed 95 percent maximum humidity (non-condensing) at any time. • Be sure to allow adequate room for proper air ventilation at the front, back, and sides of the switch. No clearance is necessary at the top or bottom of the chassis.
Electrical Requirements	OmniSwitch 9000 Series switches have the following general electrical requirements: • Each switch requires one grounded electrical outlet for each power supply installed in the chassis (up to three for OS9700 switches; up to two for OS9600 switches). OmniSwitch 9000 Series switches offer both AC and DC power supply support. For switches using AC power connections, each supplied AC power cord is 2 meters (approximately 6.5 feet) long. Do not use extension cords. **Redundant AC Power.** It is recommended that each AC outlet reside on a separate circuit. With redundant AC, if a single circuit fails, the switch's remaining power supplies (on separate circuits) are likely to remain unaffected and can, therefore, continue operating.
OmniSwitch 9600 Chassis Dimensions	
Chassis Width	19 inches
Height	9.575 inches
Height (rack units)	5.47 RU
Chassis Depth	14.432 inches

Chassis Slot Numbering

The term slot refers to the position at which a CMM or NI module is installed in chassis. CMM slot positions are designated as Slot A and Slot B. On OS9800 switches, NI slot numbers range from 1 to 16. On
OS9700 switches, NI slot numbers range from 1 to 8. On OS9600 switches, NI slot numbers range from 1 to 4.
Note. The OS9600 contains only one CMM slot.
Power supply bays are also given specific slot numbers. On OS9800 switches, power supply slot numbers are designated PS-1 through PS-4 from top to bottom. On OS9700 switches, power supply slot numbers are designated PS-1 through PS-3 from top to bottom.
On OS9600 switches, power supply slot numbers are designated PS-1 and PS-2 from top to bottom.

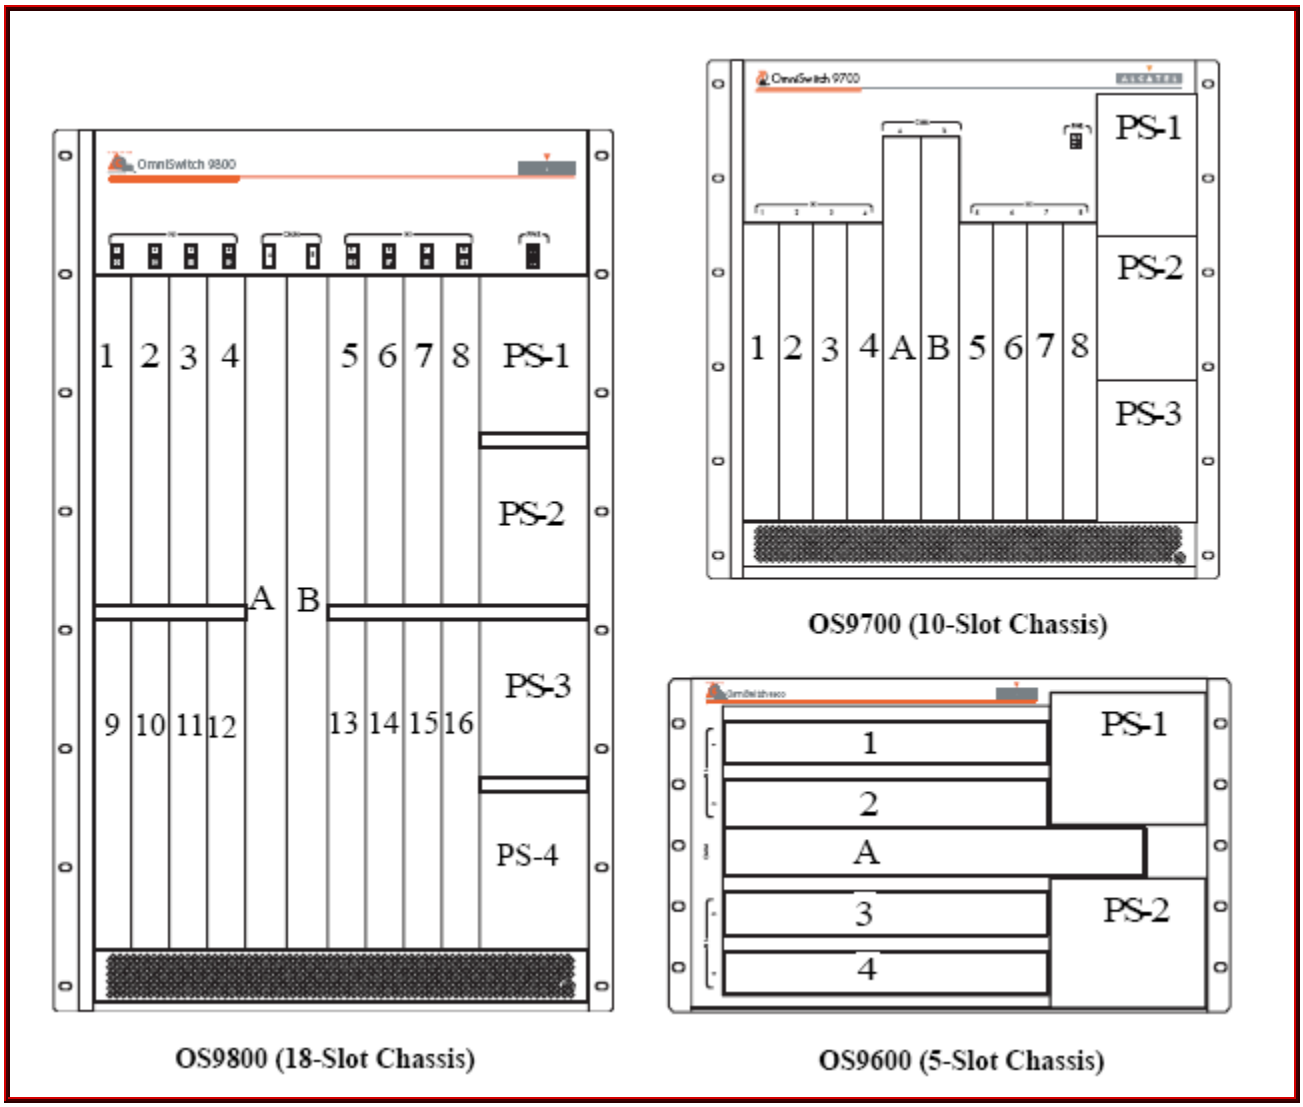

OS9000 Chassis Management Module (CMM)

The Chassis Management Module (CMM) is the management unit for OmniSwitch 9000 Series switches.
In its role as the management unit, the CMM provides key system services, including:

- Console, USB, and Ethernet management port connections to the switch
- Software and configuration management, including the Command Line Interface (CLI)
- Web-based management (WebView)
- SNMP management
- Power distribution
- Switch diagnostics
- Important availability features, including redundancy (when used in conjunction with another CMM), software rollback, temperature management, and power management
- The CMM also contains the switch fabric unit for the OmniSwitch 9000 Series switches. Data passing from one NI module to another passes through the CMM fabric. When two CMMs are installed, both fabrics are normally active.

*Note. OmniSwitch 9000 Series CMMs are colored **orange** to distinguish them from OmniSwitch 7700/7800 CMMs that are colored white. Do not install and/or mix OmniSwitch 9000 Series and OmniSwitch 7700/7800 CMMs and/or NIs in the same chassis. OmniSwitch 9000 Series CMMs will not interoperate with any of the OmniSwitch 7700/7800 CMMs/NIs in the same chassis.*

CMM Redundancy

CMM redundancy is one of the switch's most important failover features. For CMM redundancy, two fully operational CMM modules must be installed in the chassis at all times. In addition, the software on the two CMM modules must be synchronized. When two CMMs are running in the switch, one CMM has the primary role and the other has the secondary role at any given time. The primary CMM manages the current switch operations while the secondary CMM provides backup (also referred to as "failover"). In a redundant configuration, if the primary CMM fails or goes offline for any reason, the secondary CMM is instantly notified. The secondary CMM automatically assumes the primary role.
Note: CMM redundancy is not supported on the OS9600 switches.

Synchronizing the Primary and Secondary CMMs

CMM synchronization refers to the process of copying all files in the /flash/working and /flash/certified directories of the primary CMM to the /flash/working and /flash/certified directories of the secondary CMM. This ensures that the these directories match exactly on both modules, which prevents the secondary CMM from assuming the primary role with incorrect or outdated software or configuration files in the event of a primary CMM failure.
Important. In order to have effective CMM redundancy, CMM modules must be synchronized at all times.

CMM Switching Fabric

Each OS9000 CMM module contains hardware and software elements to provide management functions for the OS9000 system. The OS9000 CMM module also contains the switch fabric for the OS9000 system.
User data flowing from one NI module to another passes through the switch fabric.
The OS9000 will operate with one or two CMM modules installed.
If there are two CMM modules, one management processor is considered "primary" and is actively managing the system. The other management processor is considered "secondary" and remains ready to quickly take over management in the event of hardware or software failure on the primary. In the event of a failure, the two processors exchange roles and the secondary takes over as primary. The switch fabric on the CMM operates independently of the management processor. If there are two CMM modules installed, both fabric modules are normally active. Two CMM modules must be installed in the OS9000 to provide full fabric capacity. If there is one CMM module installed, then there is a single management feature and performance as a dual CMM system, but there is no "secondary" CMM. Hardware or software failures in the CMM will result in a system reboot. The System fabric capacity is on half of the fabric capacity of a dual CMM system.

CMM Slot Locations

In a non-redundant (i.e., single CMM) configuration, the CMM module can be installed in either Slot A or Slot B of the chassis. In a redundant CMM configuration, a CMM module is installed in both Slot A and Slot B. Both non-redundant and redundant CMM configurations can be performed on OS9700 and OS9800 switches.
Note. CMM redundancy is not supported on the OS9600 switch because it contains only one CMM slot.
Note that the CMM slots are longer than the Network Interface (NI) slots. These slots run vertically on the OS9700 and OS9800 chassis and horizontally on the OS9600 chassis. They are located near the center of the chassis.
Refer to the figures below.

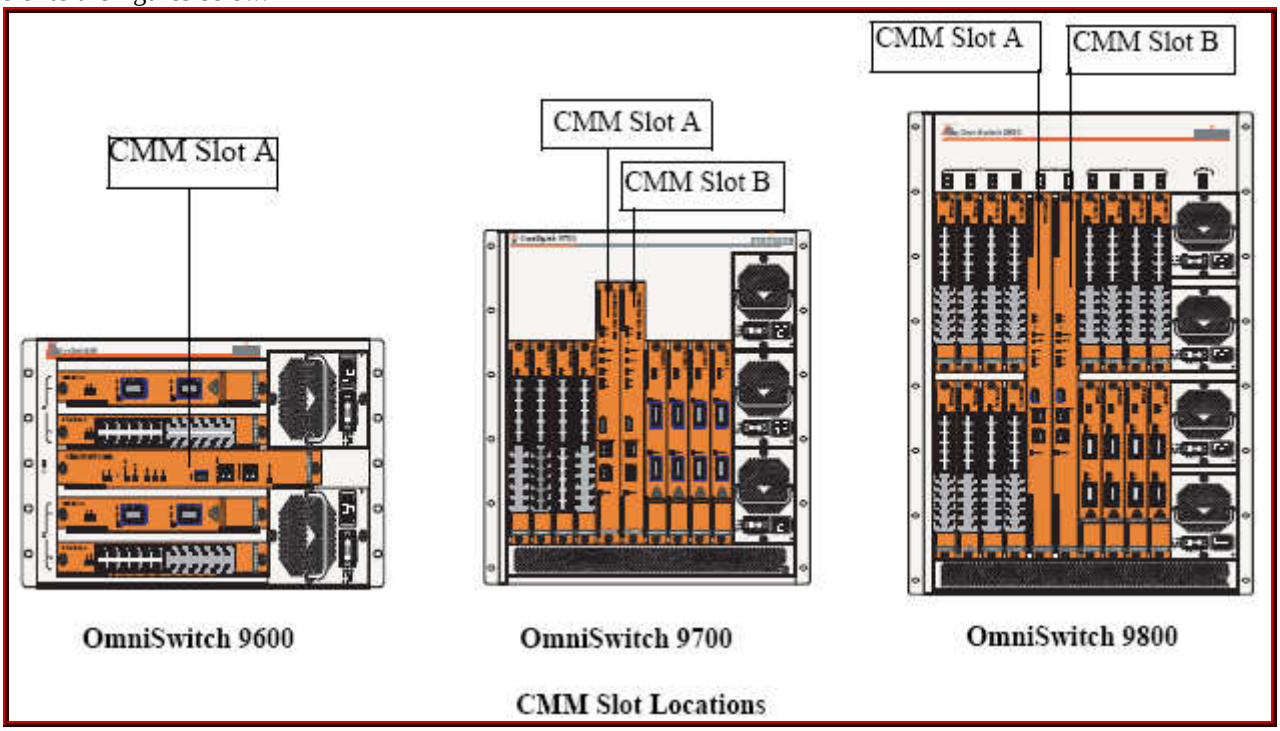

CMM Slot Locations

OS9600/OS9700-CMM Versus OS9800-CMM

OS9600/OS9700-CMM and OS9800-CMM modules offer identical functions. In addition, OS9600/OS9700-CMM and OS9800-CMM front panels provide the same port configurations and status LEDs.
However, there are two notable differences:
1 The physical dimensions of the OS9600/OS9700-CMM differ from those of the OS9800-CMM. As a result, OS9600/OS9700-CMMs and OS9800-CMMs are not interchangeable between the 9600/9700 and 9800 chassis types.
2 OS9600/OS9700-CMMs and OS9800-CMMs use identical processor boards. However, OS9800-CMMs use twice the number of network interface-related ASICs on the fabric board. This is because OmniSwitch 9800 switches support up to 16 network interface (NI) modules and OmniSwitch 9700 switches support up to 8 NI modules.

CMM Technical Specifications	
Flash Memory	**128MB**
SIMM (DRAM) Memory	**256MB**
Console port	**One RJ45 console/modem port; set to console by default**
Ethernet management port (EMP) 10/100/1000Mbps speed	**One RJ45 port; provides out-of-band network management and can be used for Telnet sessions or for downloading switch software via FTP**
USB port	**USB 2.0 supported** (the USB 2.0 will be available in a future Release)
CMM Power Consumption	**OS9600-CMM: 27watts** **OS9700-CMM: 27watts** **OS9800-CMM: 40watts**

Hot Swapping CMM Modules

Hot swapping a CMM refers to the action of adding, removing, or replacing a CMM module while the switch is operating. You are not required to enter a CLI command in order to hot swap CMM modules. This function can be performed on the fly by simply removing the module from the switch chassis.
Note. Hot Swapping the CMM module is not possible on the OS9600 because it contains only one CMM slot.

Module Presence Signaling

On-the-fly module removal is provided through the presence signaling function. All modules in the switch send out "presence signals." When a module sends out this signal, it is essentially advertising to all other modules in the switch that it is present in the chassis. When a module is present, information such as its module type (primary CMM, secondary CMM, ENI, or GNI) becomes available for monitoring functions. The presence signal is controlled through a shortened connector pin that interfaces with the switch's backplane.
Because this connector pin is shorter than the module's other backplane connectors, the presence signal connection is the first to become interrupted when a board is removed from the chassis. This allows the switch additional time (approximately 5 ms) to complete the current transfer of data before the module is completely disconnected.
(In order to avoid data loss, the switch immediately stops incoming traffic and flushes outgoing traffic on the module being removed.)
Note. Although presence signaling is designed to maintain data flow on the switch during the hot swap procedure, uninterrupted data flow cannot be guaranteed. As a result, you should not hot swap NI or CMM modules during critical network activity.

Module Types and Slot Positions

When installing modules in the chassis, consider the following:
• NI modules may be installed in any slot position from 1 through 16 in OS9800 switches, from 1 through 8 in OS9700, and from 1 through 4 in OS9600 switches.
• CMMs may be installed in slots A or B in OS9800 and OS9700, and in slot A in OS9600 switches.
• NI modules cannot be installed in CMM slots A or B; likewise, CMMs cannot be installed in any NI slot position.

Switching the Primary and Secondary Roles

The primary and secondary CMM modules can trade roles. In other words, the CMM that is currently functioning as the secondary CMM can be assigned to "take over" the role of the primary CMM. The primary CMM then assumes the secondary role. Because this action is coordinated between the two CMM modules, switch management functions are maintained during the takeover.

Chassis-Based MAC Address

The switch's base MAC address is not tied to the CMM module. Instead, the switch provides an EEPROM card near the chassis backplane that stores the MAC address. This allows the switch to retain the MAC address when a CMM module is removed or replaced.
MAC EEPROM Redundancy: A second EEPROM is provided for redundancy. An EEPROM card can be removed and replaced in the field by an authorized Alcatel.Lucent Support Engineer in the unlikely event of an EEPROM failure.

CMM Front Panel

CMM Front Panel

Module Status LEDs

OK1. Hardware Status. Displays solid green when powered on and the CMM has passed hardware diagnostic tests. Displays solid amber when powered on and the CMM has failed hardware diagnostic tests.

OK2. Software Status. Blinks green when the CMM is operational. Displays solid amber when a system software failure occurs. Blinks amber when the software is in a transitional state (e.g., when software is being downloaded to the switch).

Control/Fabrics/PSU/Temperature/Fan status LEDs

CONTROL. Displays solid green when the CMM is active, blinking green when standby, amber when malfunctioning, and blinking amber for upgrade.

FABRIC. Displays solid green when the fabric is active, blinking amber or steady amber for different fabric malfunctions.

TEMP. Displays green at 0-40°C, blinking amber at 40-45°C, and solid amber at over 45°C.

FAN. Displays solid green when all fans in the fan tray are running at normal speed. Displays solid amber if a fan error occurs (i.e., one or more fans are not running at normal speed).

PSU. Displays green when power is OK, blinking amber when one PSU is bad but the chassis has enough power, and solid amber when the chassis does not have enough power.

Ethernet Management Port LEDs

LINK. Link/Activity Status. Displays solid green when an Ethernet cable connection exists at the CMM's Ethernet Management Port. Flashes green as data is transmitted or received.

Module Status LEDs

USB Port. High speed (480 Mbps) USB 2.0 port, which can be used for quick upgrades.
Note: USB port is not supported in this release.

Console Port. The CMM's front panel provides one RJ-45 port for console connections. By default, this connector provides a DCE console connection.

Ethernet Management Port (EMP). The CMM's front panel also provides one Ethernet 10/100/1000BaseT port (copper RJ-45). This port provides out-of-band network management and can be used for Telnet sessions, switch diagnostics, and for downloading software to the switch.

This 10/100/1000BaseT port supports both 10BaseT,100BaseT and 1000BaseT with auto-negotiation through the RJ-45 connector.

Note: In redundant CMM configurations, the EMP is only operational on the *primary* CMM.

CMM Front Panel (OS9600/OS9700-CMM version shown)

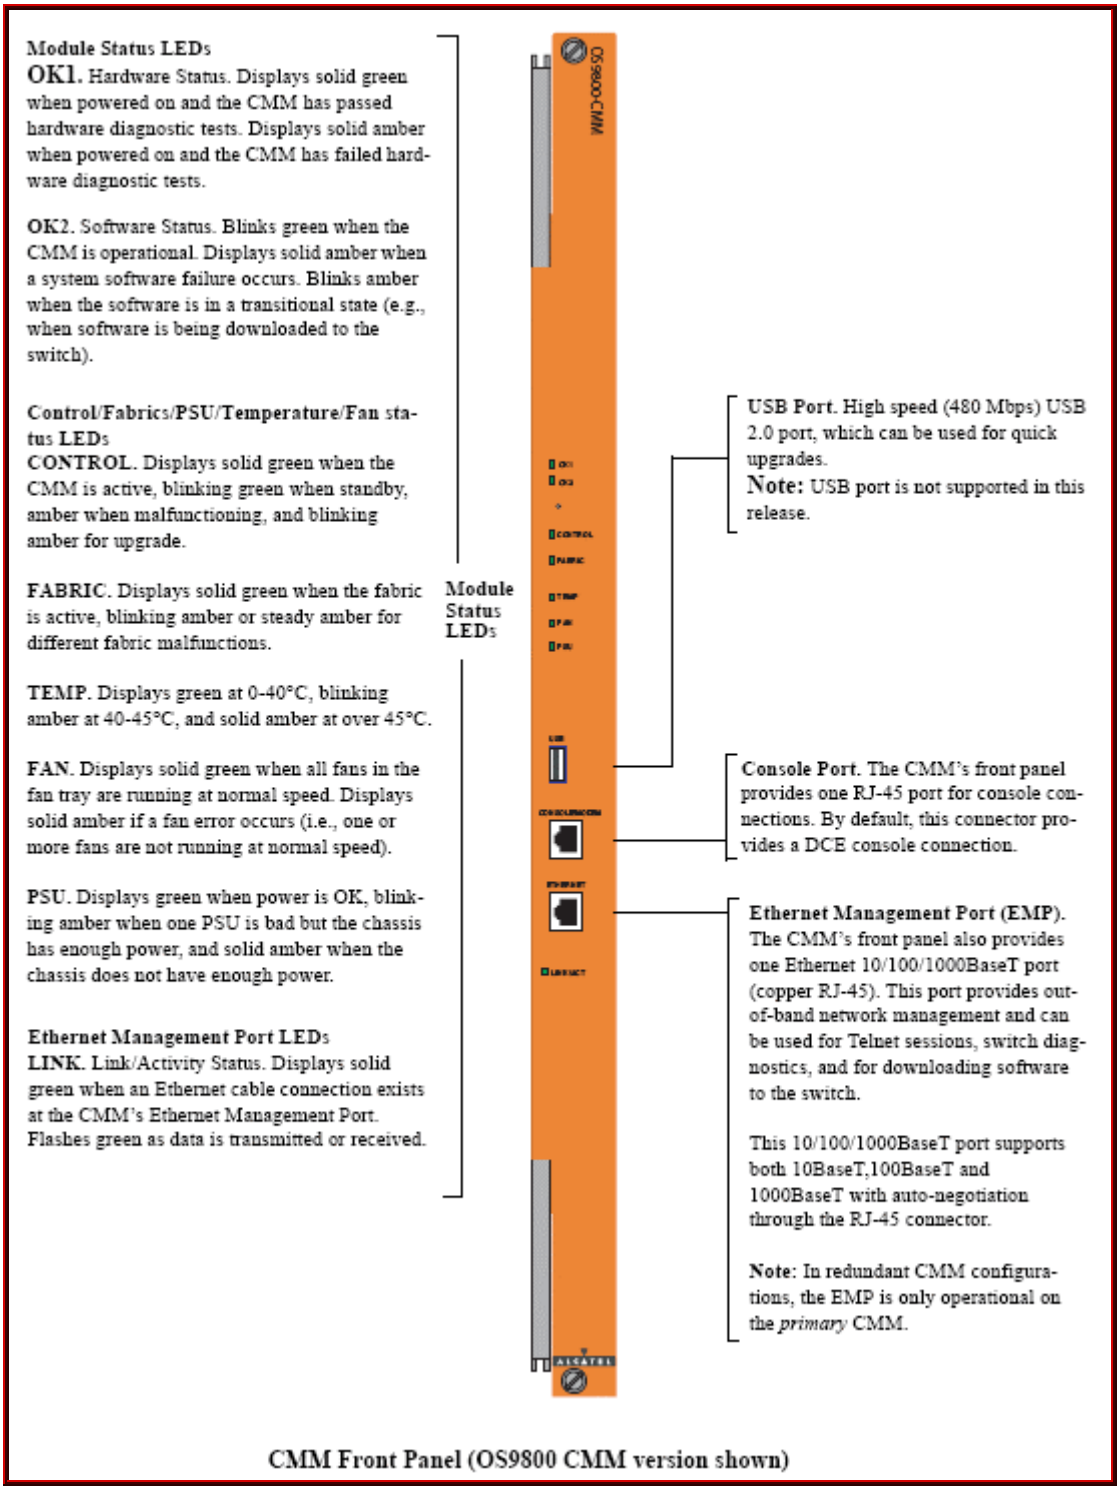

CMM Front Panel (OS9800 CMM version shown)

OS9000 Network Interface Modules

Several Gigabit Network Interface (GNI) modules and 10 Gigabit Network Interface (XNI) modules are currently available for OmniSwitch 9000 Series switches. These modules come in a variety of port speeds, including auto-sensing 10/100/1000Mbps Ethernet, Gigabit Ethernet (1Gbps), and 10 Gigabit Ethernet (10Gbps). In addition, these modules come with several connector types, including copper RJ-45 connectors on 10/100/1000 modules and LC connectors on fiber Gigabit Ethernet and 10 Gigabit Ethernet modules.

Notes.

- *OmniSwitch 9000 Series NIs is colored* **orange** *to distinguish them from OmniSwitch 7700/7800 NIs that is colored white. Do not install OmniSwitch 9000 Series and OmniSwitch 7700/7800 NIs in the same chassis.*
- *You can also manage and monitor GNI and XNI modules with WebView; Alcatel.Lucent's embedded web-based device management application. WebView is an interactive and easy-to-use GUI that can be launched from OmniVista or a web browser. Please refer to WebView's online documentation for more information.*

GNI Modules

Gigabit Ethernet Network Interface (GNI) modules provide 24 1000 Mbps (1Gbps) connections per module.
GNI modules can be used for backbone connections in networks where Gigabit Ethernet is used as the backbone media.
GNI modules can also be used in the wiring closet.
The following wire-rate GNI modules are available:

- OS9-GNI-U24. Provides 24 1000BASE-X Gigabit Ethernet SFP-MSA transceiver slots.
- OS9-GNI-C24. Provides 24 auto-sensing twisted-pair 10/100/1000BASE-T ports, individually configurable as 10BASE-T, 100BASE-TX, or 1000BASE-T.
- OS9-GNI-P24. Provides 24 auto-sensing twisted-pair Power over Ethernet (PoE) 10/100/1000BASE-T ports, individually configurable as 10BASE-T, 100BASE-TX, or 1000BASE-T.

GNI modules are supported during CMM failover.
The applicable Transceivers are hot Pluggable—i.e., they it can be installed or removed while the GNI is powered on and operating without the risk of damage to the module or the host circuitry.
When a transceiver is installed, the switch automatically gathers basic transceiver information via the connector's serial E2PROM interface. This information includes the transceiver capabilities, standard interfaces, manufacturer, and other information.
For a complete list of supported transceivers please refer to other sections in this document.
Note: Customers should use only Alcatel.Lucent-provided transceivers. Third party transceivers not provided by Alcatel.Lucent are not guaranteed to work properly.

OS9-GNI-U24 Front Panel

Module Status LEDs
OK1. Hardware Status. Displays solid green when powered on and the GNI has passed hardware diagnostic tests. Displays solid amber when powered on and the GNI has failed hardware diagnostic tests.

OK2. Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load software.

Gigabit Ethernet Port LEDs
Each fiber-based Ethernet port has a corresponding LED. This LED indicates the link and activity status for each Ethernet port. The LED displays green when a valid Ethernet cable connection exists. It flashes green as data is transmitted or received on the port.

Module Status LEDs

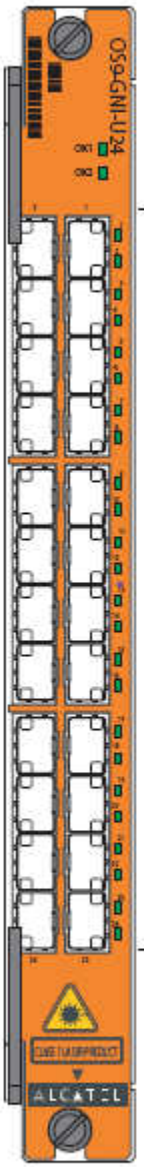

Gigabit Ethernet Ports
The OS9-GNI-U24 module provides 24 SFP transceiver slots.

OS9-GNI-U24 Front Panel

OS9-GNI-U24 Technical Specifications Overview

OS9-GNI-U24 Technical Specifications Overview	
Number of MiniGBIC ports	24 x 1000BASE- X Gigabit Ethernet SFP-MSA Hot Pluggable transceiver slots
Connector types	LC
Standards supported	IEEE 802.3z; 1000BASE-X
Data rate	1 Gigabit per second (full duplex)
Maximum frame size	9,216 bytes. OS9-GNI-U24 modules support jumbo frames (1,500 to 9,216 bytes)
MAC addresses supported	There are now two source learning modes available for the OmniSwitch 9000 Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 16K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K or more per OmniSwitch 9000 chassis. **4.1.10** The 6.1.3.R01 release provides support for this feature on the OmniSwitch 9000 Series.
Connections supported	1000BASE-X Gigabit Ethernet SFP-MSA Hot Pluggable transceiver connection supporting multiple uplinks from wire closet switches and supporting a large number of Gigabit backbone links in core applications.
Power Consumption	55watts
Gigabit Ethernet Transceivers (SFP MSA)	
SFP-GIG-47CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ gray latch. Supports single mode fiber over 1470 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-49CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ violet latch. Supports single mode fiber over 1490 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-51CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ blue latch. Supports single mode fiber over 1510 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-53CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ green latch. Supports single mode fiber over 1530 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-55CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ yellow latch. Supports single mode fiber over 1550 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-57CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ orange latch. Supports single mode fiber over 1570 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-59CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ red latch. Supports single mode fiber over 1590 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-61CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ red latch. Supports single mode fiber over 1590 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
SFP-GIG-EXTND	Extended 1000Base-SX Gigabit Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Reach of up to 2 km (based on grade and condition of fiber) on 62.5/125 µm MMF or 550m on 62.5/125 µm MMF. Requires SFP-GIG-EXTND or GBIC-GIG-EXTND at the remote termination. [Formerly known as GE-EXTND-SFP]
SFP-GIG-LH40	1000Base-LH Gigabit Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1310 nm wavelength (nominal) with an LC connector. Typical reach of 40Km on 9/125 µm SMF.
SFP-GIG-LH70	1000Base-LH Gigabit Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 70 Km on 9/125 µm SMF. [Formerly known as MINIGBIC-LH-70]
SFP-GIG-LX	1000Base-LX Gigabit Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 10Km on 9/125µm SMF. Typical reach of 550m on 50/125 & 62.5/125µm MMF. [Formerly known as MINIGBIC-LX]
SFP-GIG-SX	1000Base-SX Gigabit Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Typical reach of 300m on 62.5/125 µm MMF or 550m on 50/125 µm MMF. [Formerly known as MINIGBIC-SX]
SFP-GIG-T	1000Base-T Gigabit Ethernet Transceiver (SFP MSA) - Supports category 5, 5E, and 6 copper cabling up to 100m. SFP only works in 1000 Mbps speed and full-duplex mode
Dual Speed Ethernet Transceivers (SFP MSA)	
SFP-DUAL-MM	Dual Speed 100Base-FX or 1000Base-X Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 550m at Gigabit speed and 2km at 100Mbit speed. Notes: - At 100Mbit speed, this SFP can interoperate with SFP-100-LC-MM or similar transceiver on the other end, - At Gigabit speed, this SFP cannot interoperate with SFP-GIG-SX or similar transceiver on the other end running over 850nm wavelength. - SFP supported on OS9-GNI-U24 Gigabit Ethernet Module and OS6850-U24X SFP ports (non combo)
SFP-DUAL-SM10	Dual Speed 100Base-FX or 1000Base-X Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 10km at Gigabit speed and 100Mbit speed. Notes: - At 100Mbit speed, this SFP can interoperate with SFP-100-LC-SM15 or similar transceiver, - At Gigabit speed, this SFP can interoperate with SFP-GIG-LX or similar transceiver. - SFP supported on OS9-GNI-U24 Gigabit Ethernet Module and OS6850-U24X SFP ports (non combo)

SFP MSA Specifications

- **SFP-MSA Connector: The SFP connector consists of a 20-pin receptacle and a SFP housing cage. The 20-pin connector provides the interface for the hot Pluggable SFP module. Each SFP module contains a serial interface to provide identification information that describes the SFP capabilities, standard interfaces, manufacturer and other information.**
- **LC Connector: The LC connector is a fiber-optic cable connector that uses one-half the size of current industry standards. It increases panel density and provides duplex connection in 50% less space with duplex fits of RJ-45 footprint. It is available in SM, MM versions with Super, Ultra and Angle (APC) polishing. It provides a user-friendly audible latch to indicate proper mating and supports pull – proof.**

Notes:

- ***The worst-case Optical Power Budget in "dB" for a fiber optic link is determined by:**
  The difference between the minimum transmitters output optical power and the lowest receiver sensitivity.
- ****Maximum distance support" is claimed by the original vendor and not by Alcatel.Lucent IP Networking.**
- **Alcatel.Lucent switching & routing platforms support alternate sources of fiber-optics vendors, which are subject to change from time to time. Please be sure to contact Alcatel.Lucent Internetworking Product Marketing for a complete list of approved vendors.**
- **The following fiber optics transceivers specifications have been taken from Alcatel.Lucent IP Networking approved vendor's original Specification Sheets.**

SFP-GIG-SX Technical Specifications	
Features	**Dual data-rate of 1.25Gbps/1.0625Gbps operation** **850nm VCSEL laser and PIN photo detector** **550m transmission with 50/125µm MMF** **275m transmission with 62.5/125µm MMF** **Standard serial ID information Compatible with SFP MSA** **SFP MSA package with duplex LC connector** **With Spring-Latch for high density application** **Very low EMI and excellent ESD protection** **+3.3V single power supply** **Operating case temperature: 0 to +70°C**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match SFPs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**IEEE 802.3z, and 1000BASE-SX** **(The IEEE 802.3z standard describes the specifications for the 1000BASE-X fiber optic GigE)** **Compatible with SFP MSA** **Compatible with IEEE 802.3z** **Compatible with ANSI specifications for Fiber Channel** **Compatible with FCC 47 CFR Part 15,Class B**
Connections supported	**1000BASE-SX connections to backbone or server**
Fiber optic cable supported	**Multimode (MMF) with 62.5/125µm & 50/125µm**
Wavelength	**850nm (typical)**
Transmitter Average Output Optical Power	**Min: -9.5 dBm and Max: -4 dBm**
Receiver Sensitivity	**Min: 0 dBm and Max: -17 dBm**
Power Budget*	**7.5 dBm (17 – 9.5 = 7.5 dBm)**
Cable distances**	**≈ Supports 62.5/125µm MMF up to a maximum distance of 220 m to 300m or 50.0/125µm up to a maximum distance of 550m.**

SFP-GIG-LX Technical Specifications	
Features	**Dual data-rate of 1.25Gbps/1.0625Gbps operation** **1310nm FP laser and PIN photo detector** **550m transmission with MMF** **10km ~ 20km transmission with SMF** **Standard serial ID information compatible with SFP MSA** **SFP MSA package with duplex LC connector** **With Spring-Latch for high density application** **Very low EMI and excellent ESD protection** **+3.3V single power supply** **Operating case temperature:** **Standard: 0 to +70°C** **Industrial: -40 to +85°C**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match SFPs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**IEEE 802.3z, 1000BASE-LX** **(The IEEE 802.3z standard describes the specifications for the 1000BASE-X fiber optic Gig Eth.)** **Compatible with SFP MSA** **Compatible with IEEE 802.3z**

	Compatible with ANSI specifications for Fiber Channel Compatible with FCC 47 CFR Part 15, Class B Compatible with FDA 21 CFR 1040.10 and 1040.11, Class I Compatible with Telcordia GR-468-CORE RoHS compliance
Connections supported	1000BASE-LX connections to backbone or server
Fiber optic cable supported	Single mode (SMF) with 9/125µm Note: This transceiver also supports 50/125µm & 62.5/125µm Multimode (MMF) with a maximum distance of up to 700m (typical 550m). This special mode of operation will require "single-mode fiber offset-launch mode-conditioning patch cord" and be sure to review the IEEE 802.3 2002 clauses 38.11.1 through 38.11.4. The third party fiber optics distance extender devices will provide the capabilities of supporting much further distances than the typical 550m.
Wavelength	1310nm (typical)
Transmitter Average Output Optical Power	Min: -9.5 dBm and Max: -3 dBm
Receiver Sensitivity	Min: 0 dBm and Max: -20 dBm
Power Budget*	10.5dBm (20 – 9.5 = 10.5 dBm)
Cable distances**	≈ 10km on SMF and ≈ 700m (typical 550m) on MMF

SFP-GIG-LH70 Technical Specifications	
Features	Up to 1.25Gbps bi-directional data links 80km transmission distance with SMF 1550nm DFB laser transmitter SFP MSA package with LC optical receptacle With lever latch for high density application Single +3.3V power supply Hot-pluggable capability Low power dissipation Very low EMI and excellent ESD protection Class I laser product Monitoring interface compliant with SFF-8472 Operation case temperature: 0 to +70°C
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match SFPs on the same unit Supports operation for layer-2, and layer-3 forwarding.
Standards supported	IEEE 802.3z, 1000BASE-LH70 (The IEEE 802.3z standard describes the specifications for the 1000BASE-X fiber optic Gigabit Eth.) Compliant with SFP MSA Compliant with SFF 8472 Compliant with IEEE 802.3z Compliant with ANSI INCITS Fiber Channel FC-PI Rev13 Compliant with FCC 47 CFR Part 15, Class B Compliant with FDA 21 CFR 1040.10 and 1040.11, Class I
Connections supported	1000BASE-LH70 connections to backbone or server
Fiber optic cable supported	Single mode (SMF) with 9/125µm
Wavelength	1550nm (typical)
Transmitter Output Optical Power	Min: 0 dBm and Max: 5 dBm
Receiver Sensitivity	Min: 0 dBm and Max: -22 dBm
Power Budget*	22 dBm (22 – 0 = 22 dBm)
Cable distances**	Long reach SMF ≈ 70km

SFP-GIG-LH40 Technical Specifications	
Features	Dual data-rate of 1.25Gbps/1.0625Gbps 40km transmission distance with 9/125 µm SMF 1310nm uncooled DFB laser PIN photodiode receiver Class I laser product Digital diagnostic monitor interface Compatible with SFF-8472 SFP MSA package with duplex LC receptacle With lever latch for high density application Very low EMI and excellent ESD protection Single 3.3V power supply Operating case temperature: 0 to +70°C
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match SFPs on the same unit Supports operation for layer-2, and layer-3 forwarding
Standards supported	IEEE 802.3z, 1000BASE-LH40

	Compatible with SFP MSA Compatible with SFF-8472 Compatible with IEEE 802.3z Compatible with IEEE 802.3ah Compatible with ANSI INCITS Fiber Channel FC-PI Rev13 Compatible with FCC 47 CFR Part 15, Class B Compatible with FDA 21 CFR 1040.10 and 1040.11, Class I RoHS compliant
Connections supported	1000BASE-LH40 connections to backbone or server
Fiber optic cable supported	Single mode (SMF)
Wavelength	1310nm (typical)
Transmitter Average Output Optical Power	Min: -2 dBm and Max: 3 dBm
Receiver Sensitivity	Min: 0 dBm and Max: -22 dBm
Power Budget*	20 dBm (22 – 2 = 20 dBm)
Cable distances**	Long reach SMF ≈ 40km
SFP-GIG-EXTND Technical Specifications	
Overview	The Fiber Driver® SFP Multimode Extender increases the reach of Gigabit Ethernet and Fiber Channel data links to distances that far exceed the defined standard. This technology allows multimode (MM) fiber previously used for FDDI, Fast Ethernet and other legacy protocols to now be used for creating high-speed communication backbones. Since first appearing in the Fiber Driver Gigabit Multimode Extender module, the performance and reliability of the Multimode Extender (MMX) technology has been proven in installations throughout the world. Typically, Gigabit Ethernet and Fiber Channel transmissions distances over MM fiber are limited to 550 meters or less, far shorter than the 2-kilometer standard for when multimode fiber is used to transmit FDDI or Fast Ethernet. This fact has left IT managers needing to implement gigabit-speed protocols with little choice but to abandon their existing multimode fiber plant and install new single mode fiber.
Features	Extended 1000Base-SX Gigabit Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Reach of up to 2 km (based on grade and condition of fiber) on 62.5/125 µm MMF or 550m on 50.0/125 µm MMF. Requires SFP-GIG-EXTND or GBIC-GIG-EXTND at the remote termination. Supply Voltage: 3.3V Transmits Gigabit Ethernet (IEEE 802.3) or Fiber Channel (ANSI X3.230-1994) up to 4 km (2 km guaranteed. Maximum range depends upon grade and condition of fiber plant used) over 62.5µm and 50µm dual fiber multimode links SFP MSA SFF-8074i compliant Bellcore GR-468 compliant Multimode DSC adapter Plug-n-play, hot swappable functionality Low EMI metal enclosure Operating temperature: 0 to +70°C
Function	SFP Extended Multimode 7 SFP Extended Multimode with ROHS Compliance
Connector Type	LC
Standards supported	IEEE 802.3z
Connections supported	1000BASE-SX connections to backbone or server
Fiber optic cable supported	Multimode (MMF)
Wavelength	850nm (typical)
Transmitter Output Optical Power	Min: N/A and Max: N/A
Receiver Sensitivity	Min: N/A and Max: N/A
Power Budget*	N/A
Cable distances**	≈ 2km on 62.5/125 µm MMF or 550m on 50.0/125 µm MMF
SFP-GIG-T Technical Specifications	
Features	• Up to 1.25Gb/s bi-directional data links • Hot-pluggable SFP footprint • Extended case temperature range (0°C to +85°C) • Fully metallic enclosure for low EMI • Low power dissipation (1.2 W typical) • Compact RJ-45 connector assembly • Access to physical layer IC via 2-wire serial bus • 10/100/1000 BASE-T operation in host systems with SGMII interface
Connector Type	RJ-45
Standards supported	IEEE 802.3z, and 1000BASE-T
Connections supported	1000BASE-T connections to backbone or server
Copper Cables supported	CAT5, CAT5e, and CAT6

Cable distances	**≈ 100m at 1000Mbps and full-duplex mode**
SFP-DUAL-MM Technical Specifications	
Features	**Build-in PHY supporting SGMII Interface** **Dual data-rate of 100BASE-FX/1000BASE-LX operation** **1310nm FP laser and PIN photo-detector** **0.5m~2km transmission with MMF at 125Mbps** **0.5m~550m transmission with MMF at 1.25Gbps** **Standard serial ID information compliable with SFP MSA** **SFP MSA package with duplex LC connector** **With Spring-Latch for high density application** **Very low EMI and excellent ESD protection** **+3.3V single power supply** **Operating case temperature: 0 to +70°C**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match SFPs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**802.3z, and 100BASE-FX** **Compliable with SFP MSA** **Compliable with IEEE 802.3-2002** **Compliable with IEEE 802.3ah-2004** **Compliable with FCC 47 CFR Part 15, Class B** **Compliable with FDA 21 CFR 1040.10 and 1040.11, Class I** **Compliable with Telcordia GR-468-CORE** **RoHS compliance**
Connections supported	**1000BASE-LX or 100BASE-FX connections to backbone or server**
Fiber optic cable supported	**Multimode (MMF)**
Wavelength	**1310nm (typical)**
1000BASE-LX	
Transmitter Average Output Optical Power	**Min: -11.5 dBm and Max: -3 dBm**
Receiver Sensitivity	**Min: 0 and Max: -22 dBm**
Power Budget*	**10.5 dBm (22 – 11.5 = 10.5 dBm)**
100BASE-FX	
Transmitter Average Output Optical Power	**Min: -20.0 dBm and Max: -14.0 dBm**
Receiver Sensitivity	**Min: 0 and Max: -28 dBm**
Power Budget*	**8 dBm (28 – 20 = 8 dBm)**
Cable distances**	**550m at 1000Mbps and 2km at 100Mbps**
SFP compatibility notes	**This SFP is not supported on OmniSwitch 6850 combo ports with the exception of the OmniSwitch 6850-U24X combo ports.**
SFP-DUAL-SM10 Technical Specifications	
Features	**Build-in PHY supporting SGMII Interface** **Dual data-rate of 100BASE-LX/1000BASE-LX operation** **1310nm FP laser and PIN photo-detector** **0.5m~10km transmission with SMF** **Standard serial ID information compliable with SFP MSA** **SFP MSA package with duplex LC connector** **With Spring-Latch for high density application** **Very low EMI and excellent ESD protection** **+3.3V single power supply** **Operating case temperature: 0 to +70°C**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match SFPs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**IEEE 802.3z, 100BASE-FX, and 1000BASE-X** **Compliable with SFP MSA** **Compliable with IEEE 802.3-2002** **Compliable with IEEE 802.3ah-2004** **Compliable with FCC 47 CFR Part 15, Class B** **Compliable with FDA 21 CFR 1040.10 and 1040.11, Class I** **Compliable with Telcordia GR-468-CORE** **RoHS compliance**
Connections supported	**100BASE-LX or 1000BASE-X connections to backbone or server**
Fiber optic cable supported	**Single mode (SMF)**
Wavelength	**1310nm (typical)**
1000BASE-LX	

Transmitter Average Output Optical Power	**Min: -9.5 dBm and Max: -3 dBm**
Receiver Sensitivity	**Min: 0 and Max: -22 dBm**
Power Budget*	**12.5 dBm (22 – 9.5 = 12.5 dBm)**
100BASE-LX	
Transmitter Average Output Optical Power	**Min: -15 dBm and Max: -8 dBm**
Receiver Sensitivity	**Min: 0 and Max: -28 dBm**
Power Budget*	**13 dBm (28 – 15 = 13 dBm)**
Cable distances**	**0.5m ~ 10km transmission with SMF**
SFP compatibility notes	**This SFP is not supported on OmniSwitch 6850 combo ports with the exception of the OmniSwitch 6850-U24X combo ports.**

Coarse Wave Division Multiplexing (CWDM)

Coarse Wave Division Multiplexing (CWDM) is an optical transceiver supporting single-mode fiber over 1470nm to 1590 nm (typical) wavelengths for use with OmniSwitch 6850 Series switches. It supports IEEE 802.3z and 1000Base-LX standards. It also supports 1000Base-CWDM connection to backbone or server. CWDMs are hot-pluggable and are available for long-reach applications; the single-mode fiber cable can reach up to 62 km.

Latch Color	*Nominal Wavelength*	*Optical Link Power Budget*	*Distance*
Gray	**1471nm**	**22dBm min.**	**Up to 62km**
Violet	**1491nm**	**22dBm min.**	**Up to 62km**
Blue	**1511nm**	**22dBm min.**	**Up to 62km**
Green	**1531nm**	**22dBm min.**	**Up to 62km**
Yellow	**1551nm**	**22dBm min.**	**Up to 62km**
Orange	**1571nm**	**22dBm min.**	**Up to 62km**
Red	**1591nm**	**22dBm min.**	**Up to 62km**
Brown	**1611nm**	**22dBm min.**	**Up to 62km**

SFP-GIG-47CWD60 Technical Specifications	
Features	**Eight (8) Wavelength CWDM Transceivers** **Compliant with SFP MSA** **Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications** **Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards** **Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km** **Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825)** **Duplex LC Optical Interface** **Loss of Signal Output & TX Disable Input** **Hot-pluggable with Single +3.3V Power Supply** **CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ gray latch. Supports single mode fiber over 1471nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match CWDMs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**IEEE 802.3z, and 1000BASE-LX**
Connections supported	**1000BASE-CWDM connections to backbone or server**
Fiber optic cable supported	**Single mode (SMF) & 9/125µm**
Wavelength	**1471nm (nominal)**
Transmitter Output Optical Power	**Min: -2 dBm and Max: +3 dBm**
Input Optical Power	**Min: -24 and Max: -3 dBm**
Power Budget*	**22 dBm (24 – 2 = 22 dBm)**
Cable distances**	**Long reach single mode (SMF) ≈ 62km**

SFP-GIG-49CWD60 Technical Specifications	
Features	**Eight (8) Wavelength CWDM Transceivers** **Compliant with SFP MSA** **Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications** **Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards** **Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km** **Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825)** **Duplex LC Optical Interface** **Loss of Signal Output & TX Disable Input** **Hot-pluggable** **Single +3.3V Power Supply** **CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ violet latch. Supports single mode fiber**

	over 1491nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match CWDMs on the same unit Supports operation for layer-2, and layer-3 forwarding
Standards supported	IEEE 802.3z, and 1000BASE-LX
Connections supported	1000BASE-CWDM connections to backbone or server
Fiber optic cable supported	Single mode (SMF) & 9/125µm
Wavelength	1491nm (nominal)
Transmitter Output Optical Power	Min: -2 dBm and Max: +3 dBm
Input Optical Power	Min: -24 and Max: -3 dBm
Power Budget*	22 dBm (24 – 2 = 22 dBm)
Cable distances**	Long reach single mode (SMF) ≈ 62km
SFP-GIG-51CWD60 Technical Specifications	
Features	Eight (8) Wavelength CWDM Transceivers Compliant with SFP MSA Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825) Duplex LC Optical Interface Loss of Signal Output & TX Disable Input Hot-pluggable Single +3.3V Power Supply CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ blue latch. Supports single mode fiber over 1511nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match CWDMs on the same unit Supports operation for layer-2, and layer-3 forwarding
Standards supported	IEEE 802.3z, and 1000BASE-LX
Connections supported	1000BASE-CWDM connections to backbone or server
Fiber optic cable supported	Single mode (SMF) & 9/125µm
Wavelength	1511nm (nominal)
Transmitter Output Optical Power	Min: -2 dBm and Max: +3 dBm
Input Optical Power	Min: -24 and Max: -3 dBm
Power Budget*	22 dBm (24 – 2 = 22 dBm)
Cable distances**	Long reach single mode (SMF) ≈ 62km
SFP-GIG-53CWD60 Technical Specifications	
Features	Eight (8) Wavelength CWDM Transceivers Compliant with SFP MSA Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825) Duplex LC Optical Interface Loss of Signal Output & TX Disable Input Hot-pluggable Single +3.3V Power Supply CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ green latch. Supports single mode fiber over 1531nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match CWDMs on the same unit Supports operation for layer-2, and layer-3 forwarding
Standards supported	IEEE 802.3z, and 1000BASE-LX
Connections supported	1000BASE-CWDM connections to backbone or server
Fiber optic cable supported	Single mode (SMF) & 9/125µm
Wavelength	1531nm (nominal)
Transmitter Output Optical Power	Min: -2 dBm and Max: +3 dBm
Input Optical Power	Min: -24 and Max: -3 dBm
Power Budget*	22 dBm (24 – 2 = 22 dBm)
Cable distances**	Long reach single mode (SMF) ≈ 62km

SFP-GIG-55CWD60 Technical Specifications	
Features	**Eight (8) Wavelength CWDM Transceivers** **Compliant with SFP MSA** **Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications** **Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards** **Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km** **Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825)** **Duplex LC Optical Interface** **Loss of Signal Output & TX Disable Input** **Hot-pluggable** **Single +3.3V Power Supply** **CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ yellow latch. Supports single mode fiber over 1551nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match CWDMs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**IEEE 802.3z, and 1000BASE-LX**
Connections supported	**1000BASE-CWDM connections to backbone or server**
Fiber optic cable supported	**Single mode (SMF) & 9/125µm**
Wavelength	**1551nm (nominal)**
Transmitter Output Optical Power	**Min: -2 dBm and Max: +3 dBm**
Input Optical Power	**Min: -24 and Max: -3 dBm**
Power Budget*	**22 dBm (24 – 2 = 22 dBm)**
Cable distances**	**Long reach single mode (SMF) ≈ 62km**
SFP-GIG-57CWD60 Technical Specifications	
Features	**Eight (8) Wavelength CWDM Transceivers** **Compliant with SFP MSA** **Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications** **Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards** **Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km** **Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825)** **Duplex LC Optical Interface** **Loss of Signal Output & TX Disable Input** **Hot-pluggable** **Single +3.3V Power Supply** **CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ orange latch. Supports single mode fiber over 1571nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match CWDMs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**IEEE 802.3z, and 1000BASE-LX**
Connections supported	**1000BASE-CWDM connections to backbone or server**
Fiber optic cable supported	**Single mode (SMF) & 9/125µm**
Wavelength	**1571nm (nominal)**
Transmitter Output Optical Power	**Min: -2 dBm and Max: +3 dBm**
Input Optical Power	**Min: -24 and Max: -3 dBm**
Power Budget*	**22 dBm (24 – 2 = 22 dBm)**
Cable distances**	**Long reach single mode (SMF) ≈ 62km**
SFP-GIG-59CWD60 Technical Specifications	
Features	**Eight (8) Wavelength CWDM Transceivers** **Compliant with SFP MSA** **Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications** **Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards** **Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km** **Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825)** **Duplex LC Optical Interface** **Loss of Signal Output & TX Disable Input** **Hot-pluggable** **Single +3.3V Power Supply** **CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ red latch. Supports single mode fiber over 1591nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match CWDMs on the same unit**

	Supports operation for layer-2, and layer-3 forwarding
Standards supported	**IEEE 802.3z, and 1000BASE-LX**
Connections supported	**1000BASE-CWDM connections to backbone or server**
Fiber optic cable supported	**Single mode (SMF) & 9/125µm**
Wavelength	**1591nm (nominal)**
Transmitter Output Optical Power	**Min: -2 dBm and Max: +3 dBm**
Input Optical Power	**Min: -24 and Max: -3 dBm**
Power Budget*	**22 dBm (24 – 2 = 22 dBm)**
Cable distances**	**Long reach single mode (SMF) ≈ 62km**
SFP-GIG-61CWD60 Technical Specifications	
Features	**Eight (8) Wavelength CWDM Transceivers** **Compliant with SFP MSA** **Compatible with IEEE 802.3z Gigabit Ethernet 1000BASE-LX PMD Specifications** **Compatible with 1.062GBd Fiber Channel 100-SM-LC-L FC-PI Standards** **Minimum Optical Link Power Budgets of 22dB and 24dB to support 62km and 70km** **Eye Safe (Class I Laser Safety per FDA/CDRH & Class 1M per IEC-825)** **Duplex LC Optical Interface** **Loss of Signal Output & TX Disable Input** **Hot-pluggable** **Single +3.3V Power Supply** **CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ brown latch. Supports single mode fiber over 1611nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125 µm SMF.**
Connector Type	**The transceiver support LC connectors and are hot swappable** **Supports the ability to mix and match CWDMs on the same unit** **Supports operation for layer-2, and layer-3 forwarding**
Standards supported	**IEEE 802.3z, and 1000BASE-LX**
Connections supported	**1000BASE-CWDM connections to backbone or server**
Fiber optic cable supported	**Single mode (SMF) & 9/125µm**
Wavelength	**1611nm (nominal)**
Transmitter Output Optical Power	**Min: -2 dBm and Max: +3 dBm**
Input Optical Power	**Min: -24 and Max: -3 dBm**
Power Budget*	**22 dBm (24 – 2 = 22 dBm)**
Cable distances**	**Long reach single mode (SMF) ≈ 62km**

OS9-GNI-C24 Front Panel

Module Status LEDs

OK1. Hardware Status. Displays solid green when powered on and the GNI has passed hardware diagnostic tests. Displays solid amber when powered on and the GNI has failed hardware diagnostic tests.

OK2. Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load software.

Module Status LEDs

Ethernet Port LEDs

Each fiber-based Ethernet port has a corresponding LED. This LED indicates the link and the activity status for each Ethernet port. The LED displays green when a valid Ethernet cable connections exists. Flashes green as data is transmitted or received on the port.

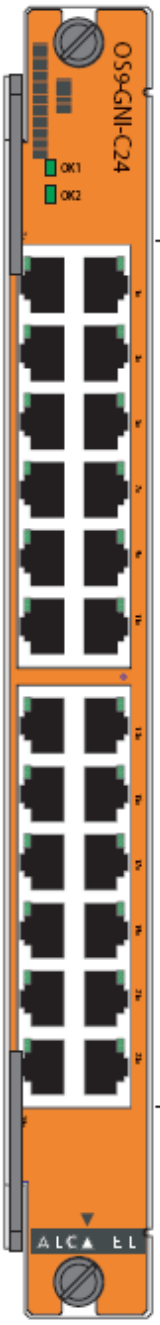

Ethernet Ports

The OS9-GNI-C24 module provide 24 10/100/1000 Ethernet ports. These ports are twisted-pair and are individually configurable as 10BaseT, 100BaseTX, or 1000BaseT. The ports use RJ-45 connectors.

OS9-GNI-C24 Technical Specifications Overview

OS9-GNI-C24 Technical Specifications Overview	
Number of MiniGBIC ports	24 x 1000BASE-T twisted Pair
Connector types	RJ-45
Standards supported	IEEE 802.3z; IEEE 802.3ab, and 1000BASE-T
Data rate	10Mbps or 100Mbps (full or half duplex) 1000Mbps (1 Gigabit per second) (full duplex)
Maximum frame size	1553 Bytes (on 10mbps or 100Mbps interfaces) 9,216 bytes (on 1 Gigabit Ethernet interfaces) OS9-GNI-C24 modules support jumbo frames (1,500 to 9,216 bytes)
MAC addresses supported	There are now two source learning modes available for the OmniSwitch 9000 Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 16K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K or more per OmniSwitch 9000 chassis. The 6.1.3.R01 release provides support for this feature on the OmniSwitch 9000 Series.
10Mbps Connections supported	10BASE-T hub or device
100Mbps Connections supported	100BASE-TX hub or device
1000Mbps Connections supported	1000BASE-T connections to backbone or server
10Mbps Cables supported	10BASE-T: unshielded twisted-pair (UTP)
100Mbps Cables supported	100BaseTX: unshielded twisted-pair (UTP), Category 5, EIA/TIA 568 or shielded twisted-pair (STP), Category-5, 100 ohm
1000Mbps Cables supported	1000BASE-T: unshielded twisted-pair (UTP), and Category 5/5e
Maximum Cable Distance	100 meters on Category 5 (any speed)
Power consumption	51watts

OS9-GNI-P24 Front Panel

Module Status LEDs

OK1. Hardware Status. Displays solid green when powered on and the GNI has passed hardware diagnostic tests. Displays solid amber when powered on and the GNI has failed hardware diagnostic tests.

OK2. Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load software.

PoE. PoE Status. This LED will be off if PoE is not available on this module and will be solid green if PoE is enabled on this module.

Module Status LEDs

Ethernet Port LEDs

Each fiber-based Ethernet port has a corresponding LED. The LED indicates the link and the activity status for each Ethernet port.

The LED displays solid green when a valid Ethernet cable connection exists and there is no PoE. Flashes green as data is transmitted or received on the port and there is no PoE.

If PoE is present, the LED displays solid amber when a valid Ethernet cable connection exists. And flashes amber as data is transmitted or received on the port if PoE is present.

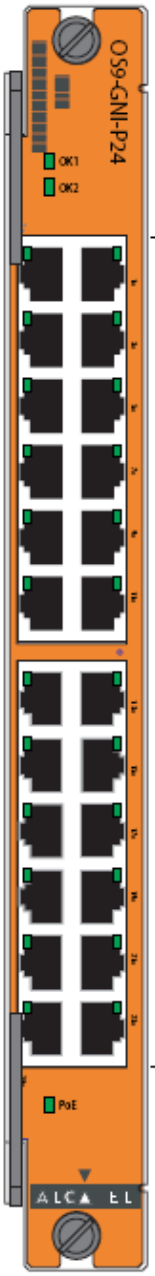

Ethernet Ports

The OS9-GNI-P24 module provides 24 10/100/1000 Power over Ethernet (PoE) ports. These ports are twisted-pair and are individually configurable as 10BaseT, 100BaseTX, or 1000BaseT. The ports use RJ-45 connectors.

Refer to the Technical Specifications table on page 5-7 for more information.

OS9-GNI-P24 Front Panel

OS9-GNI-P24 Technical Specifications Overview

OS9-GNI-P24 Technical Specifications Overview	
Number of MiniGBIC ports	24 x 1000BASE-T twisted Pair
Connector types	RJ-45
Standards supported	IEEE 802.3z; IEEE 802.3ab, 802.3af (DTE Power via MDI MIB); and 1000BASE-T
Data rate	10Mbps or 100Mbps (full or half duplex) 1000Mbps (1 Gigabit per second) (full duplex)
Maximum frame size	1553 Bytes (on 10mbps or 100Mbps interfaces) 9,216 bytes (on 1 Gigabit Ethernet interfaces) OS9-GNI-P24 modules support jumbo frames (1,500 to 9,216 bytes)
MAC addresses supported	There are now two source learning modes available for the OmniSwitch 9000 Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 16K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K or more per OmniSwitch 9000 chassis. The 6.1.3.R01 release provides support for this feature on the OmniSwitch 9000 Series.
10Mbps Connections supported	10BASE-T hub or device
100Mbps Connections supported	100BASE-TX hub or device
1000Mbps Connections supported	1000BASE-T connections to backbone or server
10Mbps Cables supported	10BASE-T: unshielded twisted-pair (UTP)
100Mbps Cables supported	100BaseTX: unshielded twisted-pair (UTP), Category 5, EIA/TIA 568 or shielded twisted-pair (STP), Category-5, 100 ohm
1000Mbps Cables supported	1000BASE-T: unshielded twisted-pair (UTP), and Category 5/5e
Maximum Cable Distance	100 meters on Category 5
Power consumption	54watts
Default amount of inline power allocated per switch slot	210 watts (please note that future enhancements will provide more PoE power per slot and/or per module)
Default amount of inline power allocated for each port	15400 Milliwatts (15.4 watts)
Range of inline power allowed for each port	3000–18000 Milliwatts (3 watts to 18 watts)
Power Over Ethernet (PoE)	IEEE 802.3af (requires OS9-GNI-P24 & PoE shelf) Maximum (assuming no P/S redundancy) power of 2100W: (4 x (600W - PSU-overhead=525W)) using the OS9-IP-SHELF (please note that 2400watts of PoE power will be supported in a future release) Maximum power of 230W / 390W using the OS9-IPS-0230A/ OS9-IPS-0390A (used exclusively on the OS9600 chassis type) The 510W and 360W (aka. 230W/390W) power supplies can be used as an alternate power source for PoE. A single 510W power supply allocates 390W for the PoE functionality; similarly, a single 360W power supply allocates 230W for the PoE functionality. Only one power supply module can be installed per switch, not both. These power modules do not support load sharing. Note. The 360W/510W power supplies are only supported on OS9600 switches and not on OS9700/OS9800 switches.

XNI Modules

OmniSwitch 9000 Series 10 Gigabit Network Interface (XNI) modules provide up to six 10000 Mbps (10Gbps) connections per module. In addition, XNI modules can be used in enterprise applications including backbone connections in networks where 10 Gigabit Ethernet is used as the backbone media. XNI modules are supported during CMM failover. The following wire-rate 10Gbps XNI modules are available:

- OS9-XNI-U2. Provides two XFP slots:

The OS9-XNI-U2 module provides two XFP slots. An XFP is a 10Gbps small-form, factor-Pluggable, module that is hot Pluggable; i.e., it can be installed or removed while the XNI is powered on and operating without the risk of damage to the XFP module or the host circuitry.

- OS9-XNI-U6. Provides six XFP slots:

The OS9-XNI-U6 module provides six XFP slots. An XFP is a 10Gbps small-form, factor-Pluggable, module that is hot Pluggable; i.e., it can be installed or removed while the XNI is powered on and operating without the risk of damage to the XFP module or the host circuitry.

OS9-XNI-U2 Technical Specifications Overview

OS9-XNI-U2 Technical Specifications Overview	
Number of XFP ports	2 x 10GBASE-X slots
Connector types	LC
Standards supported	IEEE 802.3ae & 10-Gigabit Ethernet for 10GBASE-SR/-LR/-ER/-ZR
Data rate	10 Gigabit per second (full duplex)
Maximum frame size	9,216 bytes. OS9-XNI-U2 modules support jumbo frames (1,500 to 9,216 bytes)
MAC addresses supported	There are now two source learning modes available for the OmniSwitch 9000 Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 16K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K or more per OmniSwitch 9000 chassis. The 6.1.3.R01 release provides support for this feature on the OmniSwitch 9000 Series.
Connections supported	10GBASE-S, 10GBASE-L, 10GBASE-E and 10GBASE-Z over LAN Phy.
Fiber optic cables supported	Multimode (62.5 and 50 µm) and single mode
Power Budget	XFP-10G-SR: 7.3 dB XFP-10G-LR: 9.4 dB XFP-10G-ER40: XFP-10G-ZR80:
Output optical power	XFP-10G-SR: -7.3 dBm (minimum) XFP-10G-LR: -8.2 to 0.5 dBm XFP-10G-ER40: XFP-10G-ZR80:
Input optical power	XFP-10G-SR: -9.9 to -1.0 dBm XFP-10G-LR: -14.4 to 0.5 dBm XFP-10G-ER40: XFP-10G-ZR80:
Cable Distances	XFP-10G-SR: 300 m (high modal bandwidth fiber is required to reach 300 meters) XFP-10G-LR: 10 km XFP-10G-ER40: 40 km XFP-10G-ZR80: 80 km
Power	36watts
10-Gigabit Ethernet Transceivers (XFP MSA)	
XFP-10G-ER40	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 40 Km on 9/125 µm SMF.
XFP-10G-LR	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single mode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 10 Km on 9/125 µm SMF. [Formerly known as 10G-XFP-LR]
XFP-10G-SR	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Typical reach of 300m on 50/125 µm MMF. [Formerly known as 10G-XFP-SR]
XFP-10G-ZR80	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 80 Km on 9/125 µm SMF.

OS9-XNI-U2 Front Panel

Module Status LEDs

OK1. Hardware Status. Displays solid green when powered on and the XNI has passed hardware diagnostic tests. Displays solid amber when powered on and the XNI has failed hardware diagnostic tests.

OK2. Software Status. Blinks green when the XNI is operational and has successfully loaded software. Displays solid amber when powered on and the XNI has failed to load software.

LINK/ACT LED

Each 10-gigabit port has a single LED for monitoring XFP link status and activity. The LED displays solid green when the port is up; the LED blinks green when it is transmitting or receiving packets in a link up state. The LED is off when no link is detected.

Module Status LEDs

10-gigabit Ethernet Ports

The OS9-XNI-U2 module provides 2 XFP slots.

OS9-XNI-U2 Front Panel

OS9-XNI-U6 Technical Specifications Overview

OS9-XNI-U6 Technical Specifications Overview	
Number of XFP ports	6 x 10GBASE-X slots
Connector types	LC
Standards supported	IEEE 802.3ae & 10-Gigabit Ethernet for 10GBASE-SR/-LR/-ER/-ZR
Data rate	10 Gigabit per second (full duplex)
Maximum frame size	9,216 bytes. OS9-XNI-U2 modules support jumbo frames (1,500 to 9,216 bytes)
MAC addresses supported	There are now two source learning modes available for the OmniSwitch 9000 Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 16K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K or more per OmniSwitch 9000 chassis. The 6.1.3.R01 release provides support for this feature on the OmniSwitch 9000 Series.
Connections supported	10GBASE-S, 10GBASE-L, 10GBASE-E and 10GBASE-Z over LAN Phy.
Fiber optic cables supported	Multimode (62.5 and 50 µm) and single mode
Power Budget	XFP-10G-SR: 7.3 dB XFP-10G-LR: 9.4 dB XFP-10G-ER40: XFP-10G-ZR80:
Output optical power	XFP-10G-SR: -7.3 dBm (minimum) XFP-10G-LR: -8.2 to 0.5 dBm XFP-10G-ER40: XFP-10G-ZR80:
Input optical power	XFP-10G-SR: -9.9 to -1.0 dBm XFP-10G-LR: -14.4 to 0.5 dBm XFP-10G-ER40: XFP-10G-ZR80:
Cable Distances	XFP-10G-SR: 300 m (high modal bandwidth fiber is required to reach 300 meters) XFP-10G-LR: 10 km XFP-10G-ER40: 40 km XFP-10G-ZR80: 80 km
Power	67watts
10-Gigabit Ethernet Transceivers (XFP MSA)	
XFP-10G-ER40	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 40 Km on 9/125 µm SMF.
XFP-10G-LR	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single mode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 10 Km on 9/125 µm SMF. [Formerly known as 10G-XFP-LR]
XFP-10G-SR	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Typical reach of 300m on 50/125 µm MMF. [Formerly known as 10G-XFP-SR]
XFP-10G-ZR80	10 Gigabit Ethernet optical transceiver (XFP MSA). Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 80 Km on 9/125 µm SMF.

OS9-XNI-U6 Front Pane

Module Status LEDs
OK1. Hardware Status. Displays solid green when powered on and the XNI has passed hardware diagnostic tests. Displays solid amber when powered on and the XNI has failed hardware diagnostic tests.

OK2. Software Status. Blinks green when the XNI is operational and has successfully loaded software. Displays solid amber when powered on and the XNI has failed to load software.

Module Status LEDs

LINK/ACT LED
Each 10-gigabit port has a single LED for monitoring XFP link status and activity. The LED displays solid green when the port is up; the LED blinks green when it is transmitting or receiving packets in a link up state. The LED is off when no link is detected.

10-gigabit Ethernet Ports
The OS9-XNI-U6 module provides 6 XFP slots.

OS9-XNI-U6 Front Panel

10Gbps Small Form Factor Pluggable (XFPs)

10Gbps Small Form Factor Pluggable (XFPs) are fiber-based optical transceivers. XFPs are fully hot-swappable and are available for both short-reach and long-reach applications.
The following XFP types are available:

- The XFP-10G-LR is a long-reach 10-gigabit optical transceiver that supports single mode fiber over 1310 nm wavelengths. It also supports 10 micron fiber up to a maximum distance of 10km.
- The XFP-10G-SR is a short-reach 10-gigabit optical transceiver that supports multimode fiber over 850 nm wavelengths. It also supports 50/62.5 micron fiber up to a max distance of 300m (depending on the grade of fiber used).
- The XFP-10G-ER40 is a long-reach 10-gigabit optical transceiver that supports single-mode fiber over 1550 nm wavelengths. It also supports 10 micron fiber up to a max distance of 40km.
- The XFP-10G-ZR80 is a long-reach 10-gigabit optical transceiver that supports single-mode fiber over 1550 nm wavelengths. It also supports 10 micron fiber up to a max distance of 80km (depending on the grade of fiber used).

Note: Customers should use only Alcatel.Lucent-provided XFP modules. Third party XFP modules not provided by Alcatel.Lucent are not guaranteed to work properly.

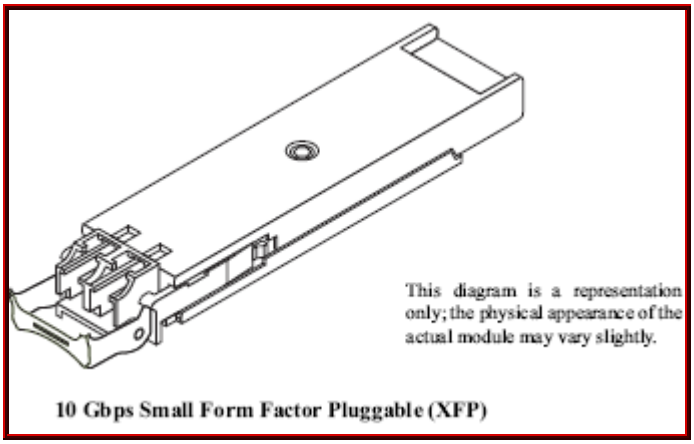

10 Gbps Small Form Factor Pluggable (XFP)

XFP-10G Specifications Eye Safety

XFP transceivers are international Class 1 laser products and are eye-safe devices when operated within the limits of manufacturers' specifications. Operating XFP transceivers in a manner inconsistent with intended usage and specification might result in hazardous radiation exposure.

XFP-10G Specifications

XFP-10G-LR Technical Specifications	
Features	■ Compact form factor according to 10 Gigabit Small Form Factor Pluggable (XFP) Multi Source Agreement, Release 3.1 ■ Multiple rate and multiple protocol support —SMF: 10GE / 10GFC / SONET / SDH —MMF: 10GE / 10GFC ■ XFP MSA compliant management and diagnostic interface ■ Z-Axis hot-plug capability ■ XFI serial data Interface via 30-Pin, XFP connector ■ Support link spans up to 10km with single mode fiber and 300m with multimode fiber ■ IEEE 802.3ae 2002 compliant —10GBASE-LR and 10GBASE-SR ■ 10GFC draft 3.0 compliant —1200-SM-LL-L and 1200-MX-SN-I ■ OC-192 SR-1/STM I64.1
Standards Supported	IEEE 802.3ae & 10GBASE-LR
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match XFPs on the same unit Supports operation for layer-2, and layer-3 forwarding
Cable Supported	Single mode; Full Duplex only
Source Type	Supports 10µm & 1310nm with a serial transceiver
Cable Distances*	≈ up to 10km
Power Consumption	≈ 2.5watts
Operating Temperature	-5 to 70°C
Transmitter Average Output (launch) optical power	Min: -5.2dBm & Max: 0.5dBm
Receiver Sensitivity	Min: -12.6dBm & Max: 0dBm
Power Budget**	7.4dBm (12.6 – 5.2 = 7.4dBm)
XFP-10G-SR Technical Specifications	
Features	■ Compact form factor according to 10 Gigabit Small Form Factor Pluggable (XFP) Multi Source Agreement, Release 3.1 ■ Multiple rate and multiple protocol support —SMF: 10GE / 10GFC / SONET / SDH —MMF: 10GE / 10GFC ■ XFP MSA compliant management and diagnostic interface ■ Z-Axis hot-plug capability ■ XFI serial data Interface via 30-Pin, XFP connector ■ Support link spans up to 10km with single mode fiber and 300m with multimode fiber ■ IEEE 802.3ae 2002 compliant —10GBASE-LR and 10GBASE-SR ■ 10GFC draft 3.0 compliant —1200-SM-LL-L and 1200-MX-SN-I ■ OC-192 SR-1/STM I64.1
Standards Supported	IEEE 802.3ae & 10GBASE-SR
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match XFPs on the same unit Supports operation for layer-2, and layer-3 forwarding
Cable Supported	Multi mode; Full Duplex only
Source Type	Supports 50/125µm & 62.5/125µm & 850nm with a serial transceiver
Cable Distances*	≈ up to 300m (based on 50/125µm MMF and a Modal bandwidth of 2000MHz*km)
Power Consumption	≈ 2.5watts
Operating Temperature	-5 to 70°C
Transmitter Average Output (launch) optical power	Min: -7.3dBm & Max: -1.0dBm
Receiver Sensitivity	Min: 0dBm & Max: -11.1dBm
Power Budget**	3.8dBm (11.1 – 7.3 = 3.8dBm)

XFP-10G-ER40 Technical Specifications	
Features	Supports 9.95Gb/s to 10.7Gb/s bit rates • Hot-pluggable XFP footprint • Maximum link length of 40km • Temperature-stabilized EML transmitter • Duplex LC connector • Power dissipation <3.5W • Built-in digital diagnostic functions • Temperature range: -5°C to 70°C
Standards Supported	IEEE 802.3ae & 10GBASE-ER
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match XFPs on the same unit Supports operation for layer-2, and layer-3 forwarding
Cable Supported	Single mode; Full Duplex only
Source Type	supports 10µm & 1550nm
Cable Distances*	≈ up to 40km
Power Consumption	≈ 3.5watts
Operating Temperature	-5 to 70°C
Transmitter Average Output (launch) optical power	Min: -1.0dBm& Max: +2.0dBm
Receiver Sensitivity @9.95Gbps to 11.1Gbps	Min: 0dBm & Max: -16dBm
Power Budget**	15dBm (16 – 1.0 = 15dBm)
XFP-10G-ZR80 Technical Specifications	
Features	Supports 9.95Gb/s to 10.7Gb/s bit rates • Hot-pluggable XFP footprint • Maximum link length of 80km • Temperature-stabilized EML transmitter • Duplex LC connector • Power dissipation <3.5W • Built-in digital diagnostic functions • Temperature range: -5°C to 70°C
Standards Supported	IEEE 802.3ae & 10GBASE-ZR
Connector Type	The transceiver support LC connectors and are hot swappable Supports the ability to mix and match XFPs on the same unit Supports operation for layer-2, and layer-3 forwarding
Cable Supported	Single Mode; Full Duplex only
Source Type	Supports 10µm & 1550nm
Cable Distances*	≈ up to 80km
Power Consumption	≈ 3.5 watts
Operating Temperature	-5 to 70°C
Transmitter Average Output (launch) optical power	Min: 0dBm & Max: +4.0dBm
Receiver Sensitivity @9.95Gbps	Min: 0dBm & Max: -24dBm
Power Budget**	24dBm (24 – 0 = 24dBm)

Notes:
***Maximum distance support" is claimed by the original vendor and not by Alcatel.Lucent IP Networking.**
****The worst-case Optical Power Budget in "dB" for a fiber optic link is determined by:**
The difference between the minimum transmitters output optical power and the lowest receiver sensitivity.

Availability Feature

The switch provides a broad variety of Availability features. Availability features are hardware and software-based safeguards that help prevent the loss of data flow in the unlikely event of a subsystem failure.
In addition, some Availability features allow you to maintain or replace hardware components without powering off your switch or interrupting switch operations. Combined, these features provide added resiliency and help ensure that your switch is consistently available for your day-to-day network operations.
Hardware-related Availability features include:
- ***Software & Hardware Redundancy***
- ***Configuration Redundancy***
- ***Link Redundancy***
- ***Smart Continuous Switching***
- ***NI Module forwarding during CMM failover***
- ***Software (Image) Rollback***
- ***Hot Swapping***
- ***Hardware Monitoring***
- ***Power Checking Sequence***

Hardware Redundancy

Hardware redundancy refers to backup hardware components. If primary hardware components fail or go offline for any reason, the redundant hardware automatically assumes the primary hardware functions (this is also referred to as failover).
The following components offer redundancy:
- ***Chassis Management Modules (CMMs)***
- ***Power Supplies***
- ***Fan Units***
- ***MAC EEPROM***

Note. Redundancy is a key Availability feature; it is recommended that you install redundant hardware components in your switch whenever possible. However, CMM redundancy is not supported on the OS9600 switch because it contains only one CMM slot.

Software Rollback

Software rollback (also referred to as image rollback) essentially allows the switch to return to a prior “last known good” version of software in the event of a system software problem. The CMM controls software rollback through its resilient directory structure design (i.e., /flash/working and /flash/certified).

Hot Swapping NI Modules

You are not required to enter a CLI command in order to hot swap NI modules. The hot swap function can be performed on the fly by simply removing the module from the switch chassis.
Hot swapping refers to the action of adding, removing, or replacing certain hardware components without powering off your switch and disrupting other components in the chassis. This feature greatly facilitates hardware upgrades and maintenance and also allows you to easily replace components in the unlikely event of hardware failure.
The following hardware components can be hot swapped:
- ***Chassis Management Modules (CMMs)***
- ***Gigabit Ethernet Network Interface modules (GNIs)***
- ***10-gigabit Ethernet Network Interface modules (XNIs)***
- ***Power supplies***
- ***Fan tray***

Hot Swapping Non-Redundant Management Modules and Power Supplies; If there is only one CMM or power supply installed in the chassis and either of these components is removed or replaced, all switch functions will stop until a replacement is installed. However, hot swapping is not possible on the OS9600 switch because it contains only one CMM slot.
Hot Swapping NI Modules; It is recommended that you hot swap NIs of the same type whenever possible. Otherwise, the network configuration may be adversely affected.

Hardware Monitoring

Automatic Monitoring

Automatic monitoring refers to the switch's inbuilt sensors that automatically monitor operations. The majority of automatic monitoring is provided by the CMM. If an error is detected (e.g., over-threshold temperature), the CMM immediately sends a trap to the user. The trap is displayed on the console in the form of a text error message. (In the case of an over-threshold temperature condition, the CMM displays an amber TEMP LED in addition to sending a trap.)

LEDs

LEDs, which provide visual status information, are provided on the CMM, NI, and power supply front panels. LEDs are used to indicate conditions, such as hardware and software status, temperature errors, link integrity, data flow, etc.

User-Driven Monitoring

User-driven hardware monitoring refers to CLI commands that are entered by the user in order to access the current status of hardware components. The user enters "show" commands that output information to the console.

Monitoring NI Modules

`Front Panel LEDs`

All NIs provide a series of status LEDs located on the front panel. These LEDs offer basic status information for the following functions:

- NI hardware operation
- NI software status
- Port link and activity status

Power Checking Sequence

The power checking sequence is another inbuilt Availability feature. This feature helps regulate power in the switch whenever the switch is booted or an NI module is installed in the chassis.
The sequence is a joint effort between the CMM, the NI modules, and the power supplies. During the boot sequence, the primary CMM automatically compares the power consumption required by installed NIs with the power available from the power supplies. If there is not adequate power to support all NIs, the CMM will power on only the supported number of NIs, starting from the first NI slot position.
Important. During the power checking sequence, CMMs receive priority and are always powered on. NI modules are then powered on sequentially by slot position. In other words, the NI in slot 1 is powered on, then slot 2, then slot 3, etc.

Module Priorities during Boot Sequence

During the power checking sequence, CMMs receive priority and are always powered on. NI modules are then powered on sequentially by slot position. In other words, the NI in slot 1 is powered on, then slot 2, then slot 3, etc.

Installing a New NI into a Running Chassis

When an NI module is installed in the chassis, only a small portion of the circuitry is initially powered up. The CMM immediately reads the incoming module's ID and determines how much power the module will require. If the number of power supplies installed in the chassis can provide sufficient power, the CMM turns on the incoming module. If the number of installed power supplies cannot provide sufficient power, the incoming NI will remain powered off.

Auto negotiation Guidelines

Please note a link will not be established on any copper Ethernet port if any one of the following is true:

- The local port advertises 100 Mbps full duplex and the remote link partner is forced to 100 Mbps full duplex.
- The local port advertises 100 Mbps full duplex and the remote link partner is forced to 100 Mbps half duplex.
- The local port advertises 10 Mbps full duplex and the remote link partner is forced to 10 Mbps full duplex.
- The local port advertises 10 Mbps full duplex and the remote link partner is forced to 10 half duplex.

This is due to the fact that when the local device is set to auto negotiating 10/100 full duplex it senses the remote device is not auto negotiating. Therefore it resolves to Parallel Detect with Highest Common Denominator (HCD), which is "10/100 Half" according to IEEE 802.3 Clause 28.2.3.1.
However, since the local device is set to auto negotiating at 10/100 full duplex it cannot form a 10/100Mbps half duplex link in any of the above mentioned cases. One solution is to configure the local device to auto negotiation, 10/100 Mbps, with auto or half duplex.

Valid Port Settings on OmniSwitch 9000 Series Switches

NI Module	*Port Number / Type*	*User-Specified Port Speed (Mbps) Supported*	*User-Specified Duplex Supported*	*Auto Negotiation Supported?*
OS9-GNI-C24	24 Copper twisted pair (RJ-45)	Auto/10/100/1000	Auto/full/half	Yes
OS9-GNI-P24	24 Copper twisted pair (RJ-45) w/PoE	Auto/10/100/1000	Auto/full/half	Yes
OS9-GNI-U24	Up to 24 high-density LC ports	1000	Full	Yes
OS9-XNI-U2	Up to 2 wire-rate fiber LC	10000	Full	Yes
OS9-XNI-U6	Up to 6 oversubscribed fiber LC	10000	Full	Yes

10/100/1000 Crossover Supported

By default, automatic crossover between MDI/MDIX (Media Dependent Interface/Media Dependent Inter-face with Crossover) media is supported on OmniSwitch 9000 Series 10/100/1000Mbps (10BASE-T/100BASE-TX/1000BASE-T) ports.
Therefore, either straight through or crossover cable can be used between two OmniSwitch 9000 Series switches as long as auto negotiation is configured on both sides of the link.

10/100 Crossover Supported

By default, automatic crossover between MDI/MDIX (Media Dependent Interface/Media Dependent Inter-face with Crossover) media is supported on OmniSwitch 9000 Series 10/100Mbps (10BASE-T/100BASE-TX) ports.
Therefore, either straight through or crossover cable can be used between two OmniSwitch 9000 Series switches as long as auto negotiation is configured on both sides of the link.

Smart Continuous Switching

In redundant CMM configurations, the switch provides support for NIs during failover. In other words, if the primary CMM fails or goes offline for any reason, NI modules will continue data transmission and routing functions during the secondary CMM's takeover process. This Availability feature is referred to as Smart Continuous Switching.
Incoming Layer 2 packets will continue to be sent to the appropriate egress port during failover.
Known routes will also be supported. (Note, however, that the NI cannot learn new routes without CMM support).
Any new route information will be ignored.) Spanning Tree will continue handling BPDUs received on the switch ports, as well as port link up and down states. The Spanning Tree topology will not be disrupted.
Note. Smart Continuous Switching is designed to maintain traffic flow only during CMM failover and is not intended to support long-term traffic flow. If both the primary and redundant CMM modules go offline or are removed from the chassis, switch operations (including all NI support) will be disabled.
However, smart continuous switching is not possible on the OS9600 switch because it contains only one CMM slot.

The OmniSwitch 9000 Series Power Supply System

The OmniSwitch 9800 switch provides space to support four load-sharing power supplies.
A fully loaded OmniSwitch 9800 switch will operate normally with three power supplies; the fourth power supply can be installed for redundancy. The OmniSwitch 9700 switch provides space to support three load-sharing power supplies.
A fully loaded OmniSwitch 9700 switch will operate normally with two power supplies; the third power supply can be installed for redundancy. The OmniSwitch 9600 switch provides space to support two load-sharing power supplies.
A fully loaded OmniSwitch 9600 switch will operate normally with one power supply; the second power supply can be installed for redundancy. The same AC-to-DC or DC-to-DC power supply is used in all chassis types. The power supplies act in a load-sharing manner and are hot swappable. Each power supply includes LEDs indicating power supply operational status ("AC OK", "DC OK", and "Over Temperature"). Each AC-to-DC or DC-to-DC power supply outputs a maximum of 24VDC with 600 watts of maximum output power, which, is converted to the required lower voltages used by all boards. Each power supply supports a protected separate power switch and a protected separate power cord. Note that, in most configurations, in addition to Power Supply Unit Redundancy, power protection can also be provided.

Chassis AC-to-DC ↑ and Chassis DC-to-DC ↓ Power Supplies

600Watt AC-to-DC Power Supply

The OmniSwitch 9800 supports a total of four power supplies; the OmniSwitch 9700 supports a total of three power supplies; the OmniSwitch 9600 supports a total of two power supplies. The power supplies are installed in the power supply bays located at the right side of the chassis.

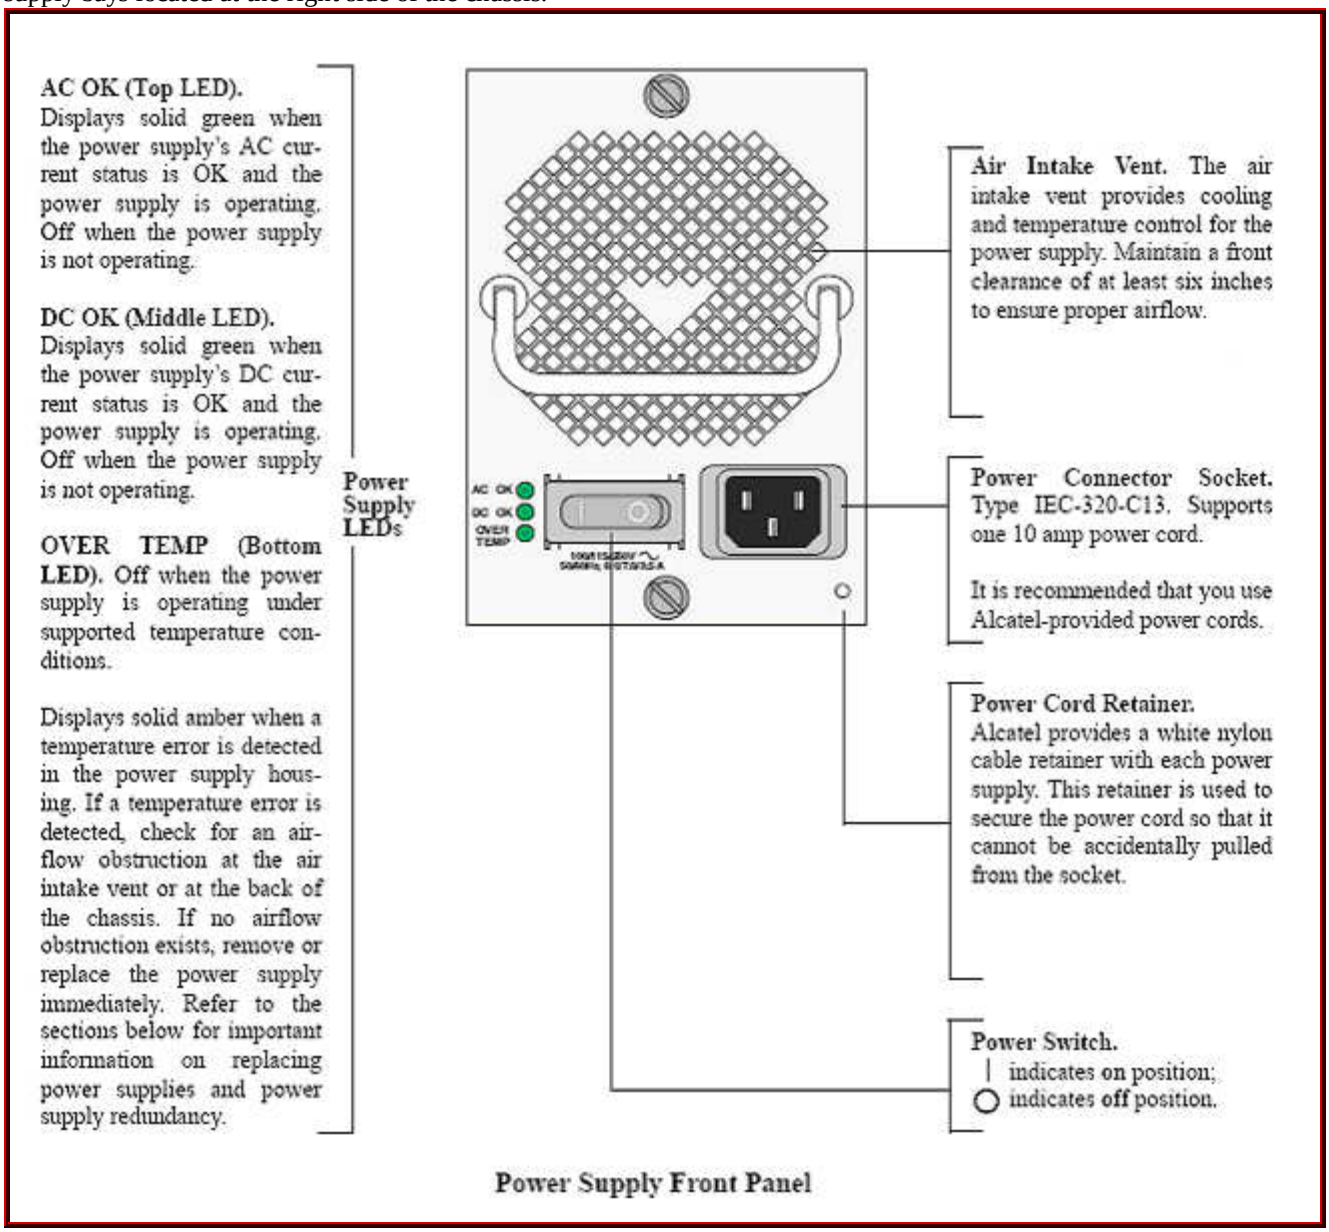

Power Supply Front Panel

600Watt DC-to-DC Power Supply

In addition to AC-to-DC power supplies, the OS9000 switches offer DC-to-DC power support (OS9-PS-0600D). The DC-to-DC power supplies are installed in the power supply bays located along the right side of the chassis.

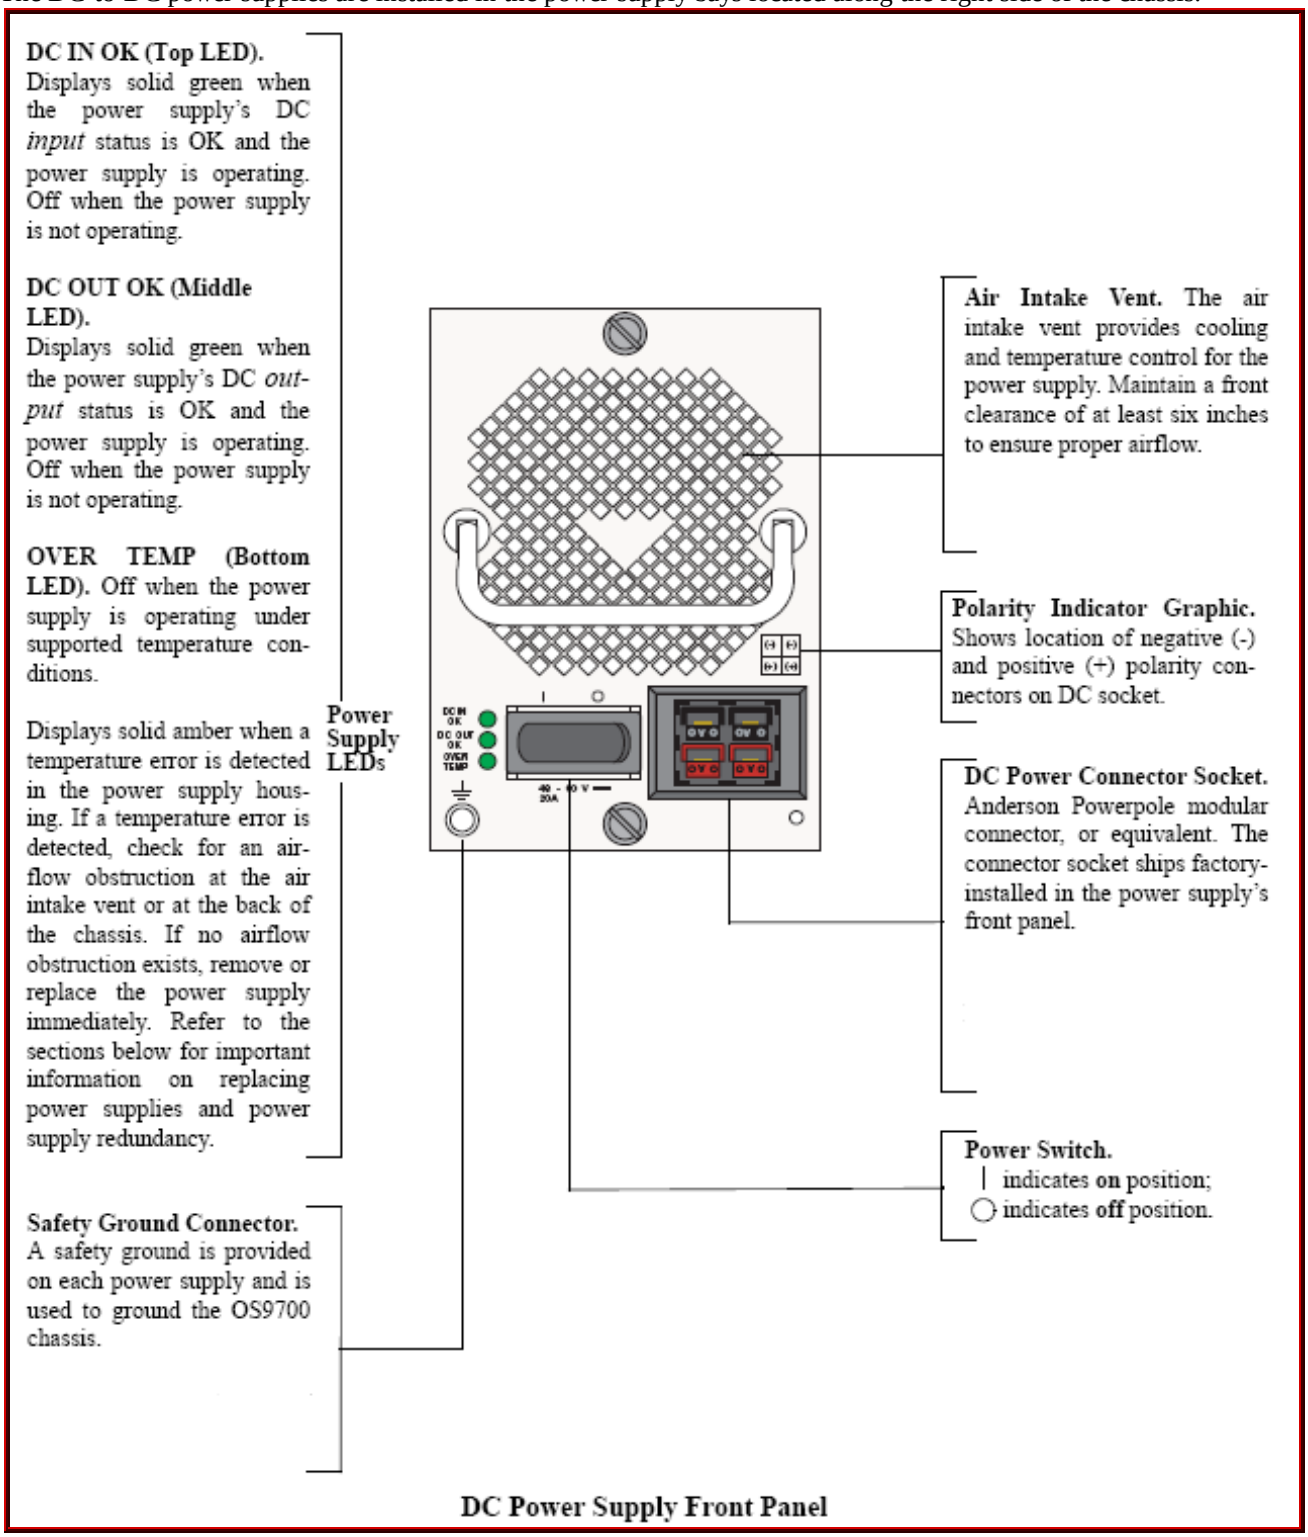

DC Power Supply Front Panel

Power Supply Specifications

Chassis AC-to-DC Power Supply

OmniSwitch 9000 Series of switches support an enclosed 600watts DC single output power supply for worldwide use as a maximum of up to N+1 power system configuration. The power supply is protected such that a short at the output to return will not result in a fire hazard, or shock hazard to the power supply. The power supply will recover automatically. The power supply design enables the removal and subsequent replacement of a defective power supply from an operating chassis, without affecting the operation of the chassis itself. Neither switching off nor removing one supply nor installing and switching on another supply causes the +24VDC line on the backplane of the chassis to vary beyond the regulation limits. This assumes that the system is configured for N+1 redundant operation. Hot swappable feature is supported.

The AC-to-DC and the DC-to-DC power supplies can be mixed and matched in the same system.

"AC-to-DC" Power Supply: OS9-PS-0600A

MTBF	BTU/hr. When the P/S is 100% loaded @ 600 watts DC	Temperature	Relative Humidity	Altitude Operating
188,000 hours	2,047.28 BTU/hr.	Operating: 0 to +70 ° C Non-Operating: – 40 to +85 ° C	Operating: 5% to 90% non-condensing, Storage: 0% to 95% non-condensing	10,000 feet @ +32°C

Rated Input Power	Rated Input Voltage	Rated Input Current	Rated Input Current	Rated Input Frequency	Efficiency
800 Watts AC	85 to 270 VAC Agency approved: 100 to 240VAC	8.0 Amps AC @ 100VAC	3.5 Amps AC @ 230VAC	47 to 63 Hz	75 % @ 115VAC +25 ° C

Rated or Maximum Output Power	Rated or Maximum Output Voltage	Rated or Maximum Output Current
600 Watts DC	24 VDC	25 Amps DC

Chassis DC-to-DC (nominal -48 VDC input) Power Supply

OmniSwitch 9000 Series of switches support an enclosed 600watts DC single output power supply for worldwide use as a maximum of up to N+1 power system configuration. The power supply is protected such that a short at the output to return does not result in a fire hazard, or shock hazard to the power supply. The power supply recovers automatically. The power supply design enables the removal and subsequent replacement of a defective power supply from an operating chassis, without affecting the operation of the chassis itself. Neither switching off nor removing one supply nor installing and switching on another supply causes the +24 VDC line on the backplane of the chassis to vary beyond the regulation limits. This assumes that the system is configured for N+1 redundant operation. The power supply supports a hot swappable feature.

The AC-to-DC and the DC-to-DC power supplies can be mixed and matched in the same system.

"DC- to-DC" Power Supply: OS9-PS-0600D

This Power Supply will have the same electrical output characteristics as that of the AC-to-DC Power Supply with the exception of the -48 VDC nominal input voltages. The use of this type of DC Power Supply is mainly intended for the Telco's/Carrier applications.

MTBF	BTU/hr. When the P/S is 100% loaded @ 600 watts DC	Temperature	Relative Humidity	Altitude Operating
188,000 hours	2,047.28 BTU/hr.	Operating: 0 to +70 °C Non-Operating: – 40 to +85 °C	Operating: 5% to 90% non-condensing, Storage: 0% to 95% non-condensing	10,000 feet @ +32°C

Rated Input Power	Rated Input Voltage	Nominal Input Voltage	Rated Input Current	Rated Input Current	Efficiency
800 Watts DC	-40 to -72VDC or Agency approved: -41 to –60VDC	-48 VDC (Minus 48VDC) The minus sign is, for polarity references only.	16.67 Amps DC @ -48VDC	20 Amps DC @ -40VDC	75 %

Rated or Maximum Output Power	Rated or Maximum Output Voltage	Rated or Maximum Output Current
600 Watts DC	24 VDC	25 Amps DC

PoE AC-to-DC Power Supply

OmniSwitch 9000 Series of switches support an enclosed 600-Watt single output AC-to-DC power supply for worldwide use as a maximum of up to N+1 power system configuration in support of PoE feature (IEEE 802.3af compliant). The power supply is protected such that a short at the output to return will not result in a fire hazard, or shock hazard to the power supply. The power supply will recover automatically. The power supply design enables the removal and subsequent replacement of a defective power supply from an operating chassis, without affecting the operation of the chassis itself. Neither switching off nor removing one supply nor installing and switching on another supply causes voltage line on the backplane of the chassis to vary beyond the regulation limits. This assumes that the system is configured for N+1 redundant operation. Hot swappable feature is supported.

"AC-to-DC" Power Supply: OS9-IPS-600A

MTBF	BTU/hr. When the P/S is 100% loaded @ 600 watts DC	Temperature	Relative Humidity	Altitude Operating
188,000 hours	2,047.28 BTU/hr.	Operating: 0 to +70 ° C Non-Operating: – 40 to +85 ° C	Operating: 5% to 90% non-condensing, Storage: 0% to 95% non-condensing	10,000 feet @ +32°C

Rated Input Power	Rated Input Voltage	Rated Input Current	Rated Input Current	Rated Input Frequency	Efficiency
800 Watts AC	85 to 270 VAC or Agency approved: 100 to 250VAC	8.0 Amps AC @ 100VAC	3.5 Amps AC @ 230VAC	47 to 63 Hz ±3%	75 % @ 115VAC +25 ° C

Rated Output Power	Rated Output Voltage	Rated Output Current
624 Watts DC	52 VDC	12 Amps DC

Maximum Output Power	Maximum Output Voltage	Maximum Output Current
600 Watts DC	52 VDC	11.5 Amps DC

OmniSwitch 9000 Series – Hardware & Software Features Overview Table

Chassis Technical Specifications	
Note: 1 inch = 2.54 centimeters & One Rack Unit = 1.75" & 1 kg = 2.2046 lbs & 1 watt ≈ 3.41214 BTU/hr.	
Rack Mountable	OmniSwitch-9800 is rack mountable in 19″ (W) and 23″ (W) racks OmniSwitch-9700 is rack mountable in 19″ (W) and 23″ (W) racks OmniSwitch-9600 is rack mountable in 19″ (W) and 23″ (W) racks **Notes:** • Due to their weight and airflow requirements, OmniSwitch 9000 Series switches cannot be wall mounted. • All OmniSwitch 9000 Series switches are shipped with integral front rack-mount flanges. These flanges support standard 19" rack mount installations. If you have non-standard Rack- mount requirements, Alcatel.Lucent offers optional hardware for the following applications: o 23" rack installations ▪ If you are installing the switch in a 23-inch wide rack, Alcatel.Lucent offers optional 23-inch rack-mounting hardware. o Side-mount hardware for additional support
Standalone	The OmniSwitch 9000 Series switches can be installed un-mounted as a standalone unit. Be sure that the installation location is a stable, flat surface that can accommodate the fully populated weight of all switches being installed. A fully populated OmniSwitch 9600 weights approximately 66 lbs (30kg); a fully populated OmniSwitch 9700 weighs approximately 128 lbs (58kg); and a fully populated OmniSwitch 9800 weights approximately 188 lbs (85kg). Note. OmniSwitch 9000 Series switches must be installed "right side up". Never attempt to operate a switch while it is lying on its side.
OmniSwitch-9800 Dimensions & Weights **Weight (fully loaded):** **< 85kg or 188lbs** **Dimensions:** **17.400″ (W) x 29.750″ (H) x 17.312″ (D)** [1]	Backplane assembly: 14.675″ (W) x 22.303″ (H) x 0.25″ (thickness) & 8 lb or 3.63 kg Power supply: 3.725″ (W) x 13.313″ (D) x 5.250″ (H) & 6 lb or 2.72 kg Fan Tray: 6.250" (W) x 17.396" (L) x 2.562" (D) & 4.5 lb or 2.04 kg (including 3 fans) Chassis assembly: 17.400″ (W) x 29.750″ (H) x 17.312″ (D) [1] & 80 lb or 36.29 kg **The OmniSwitch 9800 is 17 Rack Unit high.** NI module assembly: 9.875″ (W) x 13.024″ (D) x 1.250″ (thickness) & 3 lb or 1.36 kg Chassis Mgmt Module assembly: 21.375″ (W) x 13.024″ (D) x 1.420″ (thickness) & 8 lb or 3.63 kg **Total chassis weight fully populated per following configuration:** OS9800 Chassis including (one Chassis assembly + one Backplane assembly + one Fan-tray) + 4xP/S 2xChassis Mgmt Modules + 16xNIs (all supported NIs have the same approximate weight) 80 lb + 8 lb + 4.5 lb + 4x6 lb + 2x8 lb + 16x3 lb = **180.5 lb or 81.87 kg (± 5%)** **The shipping weight will have to include the pallet assembly plus the carton & foams weights:** Pallet assembly (to be used for shipping weight calculations only): 12 lb or 5.44 kg The carton & foams (to be used for shipping weight calculations only): approximately 10 lb or 4.53 kg Fully loaded OS9800 Chassis per above configuration + Pallet assembly + the carton & foams 180.5 lb + 12 lb + 10 lb = **202.5 lb or 91.85 kg** A typical MiniGBIC (SX or LX or LH) weighs approximately 9.07 grams **Notes:** **[1]The OmniSwitch 9800 chassis must be installed with a mandatory fan-tray assembly (OS9000-FTTC: one rear access fan-tray with three fans) for a proper switch functional operation. Therefore, the chassis depth measurement as indicated above (17.312" (D)) includes the fan tray's (2.562" (D)) depth measurement.** **1 kg = 2.2046 lbs**

OmniSwitch-9700 Dimensions & Weights **Weight (fully loaded):** **< 60kg or 133lbs** **Dimensions:** **17.400″ (W) x 19.250″ (H) x 17.312″ (D)** [1]	Backplane assembly: 14.675″ (W) x 15.022″ (H) x 0.25″ (thickness) & 7 lb or 3.17 kg Power supply: 3.725″ (W) x 13.313″ (D) x 5.250″ (H) & 6 lb or 2.72 kg Fan Tray: 6.250" (W) x 17.396" (L) x 2.562" (D) & 4.5 lb or 2.04 kg (including 3 fans) Chassis assembly: 17.400″ (W) x 19.250″ (H) x 17.312″ (D) [1] & 55 lb or 24.94 kg **The OmniSwitch 9700 is 11 Rack Unit high.** NI Module assembly: 9.875″ (W) x 13.024″ (D) x 1.250″ (thickness) & 3 lb or 1.36 kg Chassis Mgmt Module assembly: 13.083″ (W) x 13.024″ (D) x 1.420″ (thickness) & 6 lb or 2.72 kg **Total chassis weight fully populated per following configuration:** OS9700 Chassis including (one Chassis assembly + one Backplane assembly + one Fan-tray) + 3xP/S 2xChassis Mgmt Modules + 8xNIs (all supported NIs have the same approximate weight) 55 lb + 7 lb + 4.5 lb + 3x6 lb + 2x6 lb + 8x3 lb = **120.5 lb or 54.65 kg (± 5%)** **The shipping weight will have to include the pallet assembly plus the carton & foams weights:** Pallet assembly (to be used for shipping weight calculations only): 12 lb or 5.44 kg The carton & foams (to be used for shipping weight calculations only): approximately 10 lb or 4.53 kg Fully loaded OS9700 Chassis per above configuration + Pallet assembly + the carton & foams 120.5 lb + 12 lb + 10 lb = **142.5 lb or 64.63 kg** A typical MiniGBIC (SX or LX or LH) weighs approximately 9.07 grams **Notes:** **[1]The OmniSwitch 9700 chassis must be installed with a mandatory fan-tray assembly (OS9000-FTTC: one rear access fan-tray with three fans) for a proper switch functional operation. Therefore, the chassis depth measurement as indicated above (17.312" (D)) includes the fan tray's (2.562" (D)) depth measurement.** **1 kg = 2.2046 lbs**
OmniSwitch-9600 Dimensions & Weights **Weight (fully loaded):** **< 30kg or 66 lbs** **Dimensions:** **19.00″ (W) x 9.575″ (H) x 14.432″ (D)**	**Dimensions:** 19.00″ (W) x 9.575″ (H) x 14.432″ (D) **The OmniSwitch 9600 is 5.47Rack Unit high.** **Weight (fully loaded):** When fully populated (i.e., with CMM and all NI modules and power supplies installed), the OmniSwitch 9600 weighs approximately 66 lbs (30 kg).
OmniSwitch-9700 & OmniSwitch-9800 Shipping Box Dimensions	OS9700 (Half Chassis): 24″ (W) x 31″ (H) x 24″ (D) OS9800 (Full Chassis): 40″ (W) x 31″ (H) x 24″ (D) Network Interfaces (NIs): 18″ (W) x 15″ (H) x 7″ (D) OS9700-CMM: 20″ (W) x 20″ (H) x 7″ (D) OS9800-CMM: 27″ (W) x 20″ (H) x 7″ (D)
Maximum Power Consumption per board	OS9800-Chassis & Fans = 80W OS9800-CMM = 40W OS9700-Chassis & Fans = 80W OS9700-CMM = 27W OS9600-Chassis & Fans = 42W OS9600-CMM = 27W OS9-GNI-U24 = 55W (includes 24 SFP MiniGBIC Transceivers) OS9-XNI-U2 = 36W OS9-XNI-U6 = 67W OS9-GNI-C24 = 51W OS9-GNI-P24 = 54W The 1000BASE-X SFP include: SFP-GIG-SX, SFP-GIG-LX, and SFP-GIG-LH70 SFP-GIG-SX transceiver = 0.75W SFP-GIG-LX transceiver = 0.9W SFP-GIG-LH70 transceiver = 1.0W The OS9800 chassis backplane power consumption is: Negligible The OS9700 chassis backplane power consumption is: Negligible The OS9600 chassis backplane power consumption is: Negligible The PoE IP Shelf (OS9-IP-SHELF) power consumption is: Negligible The power consumption measurements were taken under fully loaded conditions. All power consumption figures include a 10% safety margin. Each main chassis AC-to-DC or DC-to-DC P/S outputs 600watts DC with an efficiency of ≥ 75%. Each PoE IP Shelf AC-to-DC P/S outputs 600watts DC with an efficiency of ≥ 75%.

Power Consumption example for OmniSwitch 9800	Power consumption calculation methodology: (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 55watts = 1040watts Total AC input power required is: 1,386.66 watts (1040 watts / 75%) **Total Power Consumption for this configuration is: 1,387watts**
Power Consumption & Heat Dissipation examples for the OmniSwitch 9800 **Note: for all practical situations, all heat dissipation calculations are based on Maximum Power Consumption (Max Power Draw)**	Power consumption calculation methodology: **Configuration Example#1:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 55watts = 1040watts Total AC input power required is: 1386.66 watts (1040 watts / 75%) 1386.66watts x 3.41214 BTU/hr. = 4,731.14 BTU/hr. **Total Heat Dissipation for this configuration is: 4,732 BTU/hr.** **Configuration Example#2:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-C24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 51watts = 976watts Total AC input power required is: 1301.33 watts (976 watts / 75%) 1301.33watts x 3.41214 BTU/hr. = 4,440.32 BTU/hr. **Total Heat Dissipation for this configuration is: 4,441 BTU/hr.** **Configuration Example#3:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-XNI-U2; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 36watts = 736watts Total AC input power required is: 981.33 watts (736 watts / 75%) 981.33watts x 3.41214 BTU/hr. = 3,348.43 BTU/hr. **Total Heat Dissipation for this configuration is: 3,349 BTU/hr.** **Configuration Example#4:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-P24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 54watts = 1,024watts Total AC input power required is: 1,365.33 watts (1,024 watts / 75%) 1365.33watts x 3.41214 BTU/hr. = 4,658.69 BTU/hr. **Total Heat Dissipation for this configuration is: 4,659 BTU/hr.** **Configuration Example#5:** **This configuration example includes a fully loaded chassis (OS9800 chassis) +** **A PoE Power Shelf (OS9-IP-SHELF) with a total of 4 x 600watts PoE P/S (OS9-IPS-0600A):** **Step#1: Calculate the heat dissipated in the main fully loaded chassis:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-P24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 54watts = 1,024watts Total AC input power required is: 1,365.33 watts (1,024 watts / 75%) 1365.33watts x 3.41214 BTU/hr. = 4,658.69 BTU/hr. Total Heat Dissipation for the fully loaded chassis: 4,659 BTU/hr. **Step#2: Calculate the heat dissipated by the 4 x 600watts PoE P/S (OS9-IPS-0600A):** Total AC input power required per PoE P/S: 800 watts (600 watts / 75%) Due to inefficiency of each PoE P/S 200watts (800watts – 600watts) is consumed per each PoE P/S: 4 x 200watts = 800watts total is consumed per a fully loaded Power Shelf 800watts x 3.41214BTU/hr. = 2,729.712 BTU/hr. Total Heat Dissipation for the fully loaded Power Shelf: 2,730 BTU/hr. **Step#3: calculate the total heat dissipation:** **Total Heat Dissipation for this configuration is: 4,659 BTU/hr. + 2,730 BTU/hr. = 7,389 BTU/hr.** **Configuration Example#6:** **This configuration example includes a OS9800-CB-A chassis + x NIs (as described) +** **A PoE Power Shelf (OS9-IP-SHELF) with a total of 3 x 600watts PoE P/S (OS9-IPS-0600A):** **Step#1: Calculate the heat dissipated in the main chassis:** (OS9800-CB-A: One OS9800 chassis + the fan-tray) + 1 x OS9800-CMM + 6 x OS9-GNI-P24 + 1 x OS9-GNI-U24 + 2 x OS9-GNI-C24; Total Sys Pwr ("load"): 80watts + 1 x 40watts + 6 x 54watts + 1 x 55watts + 2 x 51watts = 601watts Total AC input power required is: 801.33 watts (601watts / 75%) 801.33watts x 3.41214 BTU/hr. = 2,734.25 BTU/hr. Total Heat Dissipation for the fully loaded chassis: 2,735 BTU/hr. **Step#2: Calculate the heat dissipated by the 3 x 600watts PoE P/S (3 x OS9-IPS-0600A):** The PoE IP Shelf (OS9-IP-SHELF) power consumption is: Negligible Total AC input power required per PoE P/S: 800 watts (600 watts / 75%) Due to inefficiency of each PoE P/S 200watts (800watts – 600watts) is consumed per each PoE P/S: 3 x 200watts = 600watts total is consumed per a fully loaded Power Shelf 600watts x 3.41214BTU/hr. = 2,047.28 BTU/hr. Total Heat Dissipation for the fully loaded Power Shelf: 2,048 BTU/hr. **Step#3: calculate the total heat dissipation:** **Total Heat Dissipation for this configuration is: 2,735 BTU/hr. + 2,048 BTU/hr. = 4,783 BTU/hr.**

Power Consumption example for OmniSwitch 9700	Power consumption calculation methodology: (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 55watts = 574watts Total AC input power required is: 765.33 watts (574 watts / 75%) **Total Power Consumption for this configuration is: 766watts**
Power Consumption & Heat Dissipation examples for the OmniSwitch 9700 **Note: for all practical situations, all heat dissipation calculations are based on Maximum Power Consumption (Max Power Draw)**	Power consumption calculation methodology: **Configuration Example#1:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 55watts = 574watts Total AC input power required is: 765.33 watts (574 watts / 75%) 765.33watts x 3.41214 BTU/hr. = 2,611.41 BTU/hr. **Total Heat Dissipation for this configuration is: 2,612 BTU/hr.** **Configuration Example#2:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-C24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 51watts = 542watts Total AC input power required is: 722.66 watts (542 watts / 75%) 722.66watts x 3.41214 BTU/hr. = 2,465.81 BTU/hr. **Total Heat Dissipation for this configuration is: 2,466 BTU/hr.** **Configuration Example#3:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-XNI-U2; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 36watts = 422watts Total AC input power required is: 562.66 watts (422 watts / 75%) 562.66watts x 3.41214 BTU/hr. = 1,919.87 BTU/hr. **Total Heat Dissipation for this configuration is: 1,920 BTU/hr.** **Configuration Example#4:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-P24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 54watts = 566watts Total AC input power required is: 754.66 watts (566 watts / 75%) 754.66watts x 3.41214 BTU/hr. = 2,575.00 BTU/hr. **Total Heat Dissipation for this configuration is: 2,575 BTU/hr.** **Configuration Example#5:** **This configuration example includes a fully loaded chassis (OS9700 chassis) +** **A PoE Power Shelf (OS9-IP-SHELF) with a total of 4 x 600watts PoE P/S (OS9-IPS-0600A):** **Step#1: Calculate the heat dissipated in the main fully loaded chassis:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-P24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 54watts = 566watts Total AC input power required is: 754.66 watts (566 watts / 75%) 754.66watts x 3.41214 BTU/hr. = 2,575.00 BTU/hr. Total Heat Dissipation for the fully loaded chassis: 2,575 BTU/hr. **Step#2: Calculate the heat dissipated by the 4 x 600watts PoE P/S (OS9-IPS-0600A):** Total AC input power required per PoE P/S: 800 watts (600 watts / 75%) Due to inefficiency of each PoE P/S 200watts (800watts – 600watts) is consumed per each PoE P/S: 4 x 200watts = 800watts total is consumed per a fully loaded Power Shelf 800watts x 3.41214BTU/hr. = 2,729.712 BTU/hr. Total Heat Dissipation for the fully loaded Power Shelf: 2,730 BTU/hr. **Step#3: calculate the total heat dissipation:** **Total Heat Dissipation for this configuration is: 2,575 BTU/hr. + 2,730 BTU/hr. = 5,305 BTU/hr.** **Configuration Example#6:** **This configuration example includes a OS9700-CB-A chassis + x NIs (as described) +** **A PoE Power Shelf (OS9-IP-SHELF) with a total of 3 x 600watts PoE P/S (OS9-IPS-0600A):** **Step#1: Calculate the heat dissipated in the main chassis:** (OS9700-CB-A: One OS9700 chassis + the fan-tray) + 1 x OS9700-CMM + 5 x OS9-GNI-P24 + 1 x OS9-GNI-U24 + 2 x OS9-GNI-C24; Total Sys Pwr ("load"): 80watts + 1 x 27watts + 5 x 54watts + 1 x 55watts + 2 x 51watts = 534watts Total AC input power required is: 712 watts (534watts / 75%) 712watts x 3.41214 BTU/hr. = 2,429.44 BTU/hr. Total Heat Dissipation for the fully loaded chassis: 2,430 BTU/hr. **Step#2: Calculate the heat dissipated by the 3 x 600watts PoE P/S (3 x OS9-IPS-0600A):** The PoE IP Shelf (OS9-IP-SHELF) power consumption is: Negligible Total AC input power required per PoE P/S: 800 watts (600 watts / 75%) Due to inefficiency of each PoE P/S 200watts (800watts – 600watts) is consumed per each PoE P/S: 3 x 200watts = 600watts total is consumed per a fully loaded Power Shelf 600watts x 3.41214BTU/hr. = 2,047.28 BTU/hr. Total Heat Dissipation for the fully loaded Power Shelf: 2,048 BTU/hr. **Step#3: calculate the total heat dissipation:** **Total Heat Dissipation for this configuration is: 2,430 BTU/hr. + 2,048 BTU/hr. = 4,478 BTU/hr.**

Power Consumption example for OmniSwitch 9600	Power consumption calculation methodology: (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-GNI-U24; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 55watts = 289watts Total AC input power required is: 385.33 watts (289 watts / 75%) **Total Power Consumption for this configuration is: 385.33watts** To meet the required "load" for this configuration, one 600watts DC power supply is required (any additional P/S will provide load-sharing and redundancy). As required per this configuration though, Alcatel.Lucent recommends, two load-sharing P/S (includes one extra P/S for redundancy) with each one providing 144.5 watts DC output to handle the total "load".
Heat Dissipation example for the OmniSwitch 9600 **Note: for all practical situations, all heat dissipation calculations are based on Maximum Power Consumption (Max Power Draw)**	Power consumption calculation methodology: **Configuration Example#1:** (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-GNI-U24; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 55watts = 289watts Total AC input power required is: 385.33 watts (289 watts / 75%) 385.33watts x 3.41214 BTU/hr. = 1,314.79BTU/hr. **Total Heat Dissipation for this configuration is: 1,315 BTU/hr.** **Configuration Example#2:** (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-GNI-C24; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 51watts = 273watts Total AC input power required is: 364 watts (273 watts / 75%) 364watts x 3.41214 BTU/hr. = 1,242 BTU/hr. **Total Heat Dissipation for this configuration is: 1,242 BTU/hr.** **Configuration Example#3:** (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-XNI-U2; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 36watts = 213watts Total AC input power required is: 284 watts (213 watts / 75%) 284watts x 3.41214 BTU/hr. = 969.04 BTU/hr. **Total Heat Dissipation for this configuration is: 969 BTU/hr.** **Configuration Example#4:** (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-XNI-U6; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 67watts = 337watts Total AC input power required is: 449 watts (337 watts / 75%) 449watts x 3.41214 BTU/hr. = 1,533.18 BTU/hr. **Total Heat Dissipation for this configuration is: 1,533 BTU/hr.**
Power Supply Requirements	
Chassis Power Supply Requirements **OS9-PS-0600A & OS9-PS-0600D** **Power Supply Efficiency ≥ 75%**	**OmniSwich-9800:** The chassis accommodates up to four 600 watts DC maximum output power supplies in a N+1 redundancy configuration or up to three 600 watts DC maximum output power supplies in a non-redundant configuration. The required number of power supplies per chassis is dependent on the chassis configuration & load. **OmniSwitch-9700:** The chassis accommodates up to three 600 watts DC maximum output power supplies in a N+1 redundancy configuration or up to two 600 watts DC maximum output power supplies in a non-redundant configuration. The required number of power supplies per chassis is dependent on the chassis configuration & load. **OmniSwitch-9600:** The chassis accommodates up to two 600 watts DC maximum output power supplies in a N+1 redundancy configuration or one 600 watts DC maximum output power supplies in a non-redundant configuration. The required number of power supplies per chassis is dependent on the chassis configuration & load. **OS9-PS-0600D: The OmniSwitch 9800, the OmniSwitch 9700, and the OmniSwitch 9600 support a DC version Power Supply with a nominal -48 VDC input power.** **The AC-to-DC and the DC-to-DC power supplies can be mixed and matched in the same system.** **Notes:** **Refer to the section on "Power Supply Requirements per Chassis Configuration".** Each chassis power supply provides a protected power switch and a separate power cord All chassis power supplies operate in a load sharing, auto ranging & auto-sensing mode for the worldwide use. Each chassis power supply is hot swappable and occupies one P/S slot Chassis power supplies are interchangeable between the two chassis type

AC-to-DC Power Supply Input & Output Electrical Parameters

Chassis AC-to-DC Power Supply Input Parameters **OS9-PS-0600A** **This power supply is common between OS9600, OS9700 chassis and the OS9800 chassis**	**OS9-PS-0600A single AC-to-DC power supply rated input electrical parameters:** Input Power: 800 watts (600 watts / 75%) Input Voltage: 85 to 270VAC auto-ranging (Agency approved unit is indicated as 100 to 240VAC) P/S rating as indicated on the unit: 100VAC input voltage @ 8.0 Amps AC input current 110VAC input voltage @ 7.5 Amps AC input current 115VAC input voltage @ 7.0 Amps AC input current 220VAC input voltage @ 3.5 Amps AC input current 230VAC input voltage @ 3.5 Amps AC input current Input Frequency: 47 to 63 (± 3%) HZ Efficiency ≥ 75% Power Factor: The power factor, when measured over the input range of 90 to 240 VAC maximum load, shall be at least 0.95 within the operating temperature (0 to +70 °C) Note: For electrical circuit breaker design, the P/S rated electrical parameters must be considered. **Input electrical parameters for the following OmniSwitch 9800 configuration example:** **Configuration Example#1:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 55watts = 1040watts Total AC input power required is: 1386.66 watts (1040 watts / 75% efficiency factor) To meet the required power consumption for the above configuration, three 600 watts DC output power supplies are recommended (two P/S is required based on the "load") and since the power supplies Load-share, each one will provide 462.22 watts AC input for a total input power of 1386.66watts DC: [1040 watts DC output / 75 % efficiency factor]. Input Power: 1386.66watts AC Nominal Input Voltage: USA: 110VAC or Europe: 220VAC Input Current: 12.60 Amps AC @110VAC and 6.30 Amps AC @220VAC **Configuration Example#2:** **This configuration example includes a OS9800-CB-A chassis + x NIs (as described) +** **A PoE Power Shelf (OS9-IP-SHELF) with a total of 3 x 600watts PoE P/S (OS9-IPS-0600A):** **Step#1: Calculate the input parameters for the main chassis:** (OS9800-CB-A: One OS9800 chassis + the fan-tray) + 1 x OS9800-CMM + 6 x OS9-GNI-P24 + 1 x OS9-GNI-U24 + 2 x OS9-GNI-C24; Total Sys Pwr ("load"): 80watts + 1 x 40watts + 6 x 54watts + 1 x 55watts + 2 x 51watts = 601watts Total AC input power required is: 801.33 watts (601watts / 75%) Input Power: 801.33watts AC Nominal Input Voltage: USA: 110VAC or Europe: 220VAC Input Current: 7.28 Amps AC @110VAC and 3.64 Amps AC @220VAC **Step#2: Calculate the input parameters for the 3 x 600watts PoE P/S (3 x OS9-IPS-0600A):** The PoE IP Shelf (OS9-IP-SHELF) power consumption is: Negligible Total AC input power required per PoE P/S: 800 watts (600 watts / 75%) Input Voltage: 85 to 270VAC auto-ranging (Agency approved unit is indicated as 100 to 240VAC) P/S rating as indicated on the unit: 100VAC input voltage @ 8.0 Amps AC input current per P/S (3xP/S: 3 x 8.0 Amps = 24.0 Amps AC) 110VAC input voltage @ 7.5 Amps AC input current per P/S (3xP/S: 3 x 7.5 Amps = 22.5 Amps AC) 115VAC input voltage @ 7.0 Amps AC input current per P/S (3xP/S: 3 x 7.0 Amps = 21.0 Amps AC) 220VAC input voltage @ 3.5 Amps AC input current per P/S (3xP/S: 3 x 3.5 Amps = 10.5 Amps AC) 230VAC input voltage @ 3.5 Amps AC input current per P/S (3xP/S: 3 x 3.5 Amps = 10.5 Amps AC) Input Frequency: 47 to 63 (± 3%) HZ Efficiency ≥ 75% **Input electrical parameters for the following OmniSwitch 9700 configuration example:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 55watts = 574watts Total AC input power required is: 765.33 watts (574 watts / 75% efficiency factor) To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load") and since the power supplies Load-share, each one will provide 382.66 watts AC input for a total input power of 765.33 watts DC: [574 watts DC output / 75 % efficiency factor]. Input Power: 765.33watts AC Nominal Input Voltage: USA: 110VAC or Europe: 220VAC

	Input Current: 6.96 Amps AC @110VAC and 3.48 Amps AC @220VAC **Input electrical parameters for the following OmniSwitch 9600 configuration example:** (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-GNI-U24; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 55watts = 289watts Total AC input power required is: 385.33 watts (289 watts / 75% efficiency factor) To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load") and since the power supplies Load-share, each one will provide 192.66 watts AC input for a total input power of 385.33 watts DC: [289 watts DC output / 75 % efficiency factor]. Input Power: 385.33watts AC Nominal Input Voltage: USA: 110VAC or Europe: 220VAC Input Current: 3.5 Amps AC @110VAC and 1.75 Amps AC @220VAC
Chassis AC-to-DC Power Supply Output Parameters **OS9-PS-0600A**	**OS9-PS-0600A single AC-to-DC power supply rated or maximum output electrical parameters:** Output Power: 600 watts DC, Output Voltage: 24 Volts DC, Output Current: 25 Amps DC **Output electrical parameters for the following OmniSwitch-9800 configuration example:** **Configuration Example#1:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 55watts = 1040watts Output Power: 1040watts DC Output Voltage: 24VDC Output Current: 43.33 Amps DC To meet the required power consumption for the above configuration, three 600 watts DC output power supplies are recommended (two P/S is required based on the "load"). The Power Supplies load-share, each one will provide 346.66 watts DC output. Note: The 24VDC is broken down on a DC-to-DC converter (located on every module) to: 3.3VDC, 2.5VDC, and 1.8VDC required voltages **Configuration Example#2:** **This configuration example includes a OS9800-CB-A chassis + x NIs (as described) +** **A PoE Power Shelf (OS9-IP-SHELF) with a total of 3 x 600watts PoE P/S (OS9-IPS-0600A):** **Step#1: Calculate the output parameters for the main chassis:** (OS9800-CB-A: One OS9800 chassis + the fan-tray) + 1 x OS9800-CMM + 6 x OS9-GNI-P24 + 1 x OS9-GNI-U24 + 2 x OS9-GNI-C24; Total Sys Pwr ("load"): 80watts + 1 x 40watts + 6 x 54watts + 1 x 55watts + 2 x 51watts = 601watts Output Power: 601watts DC Output Voltage: 24VDC Output Current: 25.04 Amps DC **Step#2: Calculate the output parameters for the 3 x 600watts PoE P/S (3 x OS9-IPS-0600A):** Per PoE P/S: Output Power: 600 watts DC, Output Voltage: 24 Volts DC, Output Current: 25 Amps DC **Output electrical parameters for the following OmniSwitch-9700 configuration example:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 55watts = 574watts Output Power: 574watts DC Output Voltage: 24VDC Output Current: 23.91 Amps DC To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load"). The Power Supplies load-share, each one will provide 287.0 watts DC output. Note: The 24VDC is broken down on a DC-to-DC converter (located on every module) to: 3.3VDC, 2.5VDC, and 1.8VDC required voltages **Output electrical parameters for the following OmniSwitch-9600 configuration example:** (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-GNI-U24; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 55watts = 289watts Output Power: 289watts DC Output Voltage: 24VDC Output Current: 12.04 Amps DC To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load"). The Power Supplies load-share, each one will provide 144.5 watts DC output. Note: The 24VDC is broken down on a DC-to-DC converter (located on every module) to: 3.3VDC, 2.5VDC, and 1.8VDC required voltages

DC-to-DC Power Supply Input & Output Electrical Parameters	
Chassis DC-to-DC Power Supply Input Parameters **OS9-PS-0600D** **This power supply is common to OS9600, OS9700 and the OS9800 chassis**	**OS9-PS-0600D single DC-to-DC power supply rated input electrical parameters:** Input Power: 800 watts Input Voltage Range: -40 to –72 VDC or Agency Approved: -41 to –60VDC Nominal Input Voltage: -48VDC (Minus 48 VDC. The minus sign is for polarity references only) Input Current: 16.67Amps DC @ -48VDC or 20 Amps DC @ -40VDC Input Current: 13.33 Amps DC @ -60VDC or 19.5 Amps DC @ -41VDC P/S Efficiency ≥ 75% **Input electrical parameters for the following OmniSwitch 9800 configuration example:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 55watts = 1040watts Total AC input power required is: 1386.66 watts (1040 watts / 75% efficiency factor) To meet the required power consumption for the above configuration, three 600 watts DC output power supplies are recommended (two P/S is required based on the "load") and since the power supplies load-share, each one will provide 462.22watts AC input for a total input power of 1386.66watts DC: [1040watts DC output / 75 % efficiency factor]. Input Power: 1386.66watts Nominal Input Voltage: -48VDC Input Current: 28.88 Amps DC @-48VDC **Input electrical parameters for the following OmniSwitch 9700 configuration example:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 55watts = 574watts Total AC input power required is: 765.33 watts (574 watts / 75% efficiency factor) To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load") and since the power supplies Load-share, each one will provide 382.66 watts AC input for a total input power of 765.33 watts DC: [574watts DC output / 75 % efficiency factor]. Input Power: 765.33watts Nominal Input Voltage: -48VDC Input Current: 15.94 Amps DC @-48VDC **Input electrical parameters for the following OmniSwitch 9600 configuration example:** (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-GNI-U24; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 55watts = 289watts Total AC input power required is: 385.33 watts (289 watts / 75% efficiency factor) To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load") and since the power supplies Load-share, each one will provide 192.66 watts AC input for a total input power of 385.33 watts DC: [289watts DC output / 75 % efficiency factor]. Input Power: 385.33watts Nominal Input Voltage: -48VDC Input Current: 8.02 Amps DC @-48VDC
Chassis DC-to-DC Power Supply Output Parameters **OS9-PS-0600D**	**OS9-PS-0600D single DC-to-DC power supply rated or maximum output electrical parameters:** Output Power: 600 watts DC, Output Voltage: 24VDC, Output Current: 25 Amps DC **Output electrical parameters for the following OmniSwitch-9800 configuration example:** (One OS9800 chassis + the fan-tray) + 2 x OS9800-CMM + 16 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 40watts + 16 x 55watts = 1040watts Output Power: 1040watts DC Output Voltage: 24VDC Output Current: 43.33 Amps DC To meet the required power consumption for the above configuration, three 600 watts DC output power supplies are recommended (two P/S is required based on the "load"). The Power Supplies load-share, each one will provide 346.66 watts DC output. Note: The 24VDC is broken down on a DC-to-DC converter (located on every module) to: 3.3VDC, 2.5VDC, and 1.8VDC required voltages **Output electrical parameters for the following OmniSwitch-9700 configuration example:** (One OS9700 chassis + the fan-tray) + 2 x OS9700-CMM + 8 x OS9-GNI-U24; Total System Power ("load"): 80watts + 2 x 27watts + 8 x 55watts = 574watts Output Power: 574watts DC Output Voltage: 24VDC Output Current: 23.91 Amps DC To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load"). The Power Supplies load-share, each one will provide 287.70 watts DC output. Note: The 24VDC is broken down on a DC-to-DC converter (located on every module) to: 3.3VDC, 2.5VDC, and 1.8VDC required voltages

	Output electrical parameters for the following OmniSwitch-9600 configuration example: (One OS9600 chassis + the fan-tray) + 1 x OS9600-CMM + 4 x OS9-GNI-U24; Total System Power ("load"): 42watts + 1 x 27watts + 4 x 55watts = 289watts Output Power: 289watts DC Output Voltage: 24VDC Output Current: 12.04 Amps DC To meet the required power consumption for the above configuration, two 600 watts DC output power supplies are recommended (one P/S is required based on the "load"). The Power Supplies load-share, each one will provide 144.5 watts DC output. Note: The 24VDC is broken down on a DC-to-DC converter (located on every module) to: 3.3VDC, 2.5VDC, and 1.8VDC required voltages
Hardware Technical Specifications	
The OmniSwitch 9000 Series platforms include, the OS9600 Model, OS9700 Model, and OS9800 The OmniSwitch 9000 Series is a powerful layer-2& layer-3 device with wire speed switching, routing and QoS coupled with unique redundancy features that set it apart from other layer-2 & layer-3 switches on the market. The OS9000 Series runs the AOS software making it completely compatible with the OS6600 Series, OS6800 Series, OS6850 Series, OS7000 Series and OS8000 switches. When placed in the proper environment, the OS9000 Series is a very powerful and effective core switch. The OmniSwitch 9000 Series benefits from a distributed switch architecture that provides redundancy of critical hardware and software elements for a continuous traffic processing in any network conditions **without** a single point of failure. **OmniSwitch 9000 Series Switch Processing Scheme; Non-blocking, and store-and-forward**	
Chassis options	**OmniSwitch-9800: 18-slot chassis** 16 slots for NI modules + 2 slots for the management modules (OS9800-CMM) One CMM is required, one extra & optional CMM, **Backplane capacity: 1.92Tbps max.** **Switching Capacity:** • 768Gbps max with dual CMMs • 384Gbps max with single CMM **Throughput 571.4Mpps** **OmniSwitch-9700: 10-slot chassis** 8 slots for NI modules + 2 slots for the management modules (OS9700-CMM) One CMM is required, one extra & optional CMM, **Backplane capacity: 960Gbps max.** **Switching Capacity:** • 384Gbps max with dual CMMs • 192Gbps max with single CMM **4.2.3** **Throughput 285.7Mpps** **OmniSwitch-9600: 5-slot chassis** 4 slots for NI modules + 1 slot for the management modules (OS9600-CMM) One CMM is required. There is no CMM redundancy. **Backplane capacity: 960Gbps max.** **Switching Capacity:** • 192Gbps max with single CMM **Throughput 142.85 Mpps** The network interface modules (including the MiniGBICs), the power supplies and the fan tray is interchangeable between the various chassis types. The only module that is not interchangeable between the various chassis options is the Chassis Management Module (OS9800-CMM & OS9700-CMM & OS9600-CMM). The management modules provide the same functionality for all chassis types but are offered with different physical sizes. Please note that the OS9600-CMM and the OS9700-CMM are of the same physical size.
The OmniSwitch 9800 chassis Note. OmniSwitch 9800 NI modules and OmniSwitch 7000 NI modules should not be mixed in the same Chassis.	The OmniSwitch 9800 chassis supports a high-performance switch fabric and provides 16 slots for Ethernet, Gigabit Ethernet, and/or 10 Gigabit Ethernet Network Interface (NI) modules. An additional two slots are reserved for primary and redundant Chassis Management Modules (CMMs). The OmniSwitch 9800 supports a maximum of four power supplies and up to 384 10/100/1000 copper ports and/or 1000 Mbps fiber ports. It is suitable for wiring closet installations. It can also be equipped with up to 96 10 Gigabit Ethernet ports for use as the core switch.
The OmniSwitch 9700 chassis Note. OmniSwitch 9700 NI modules and OmniSwitch 7000 NI modules should not be mixed in the same Chassis.	The OmniSwitch 9700 chassis supports a high-performance switch fabric and provides 8 slots for Ethernet, Gigabit Ethernet, and/or 10 Gigabit Ethernet Network Interface (NI) modules. An additional two slots are reserved for primary and redundant Chassis Management Modules (CMMs). The OmniSwitch 9700 supports a maximum of three power supplies and up to 192 10/100/1000 copper ports and/or 1000 Mbps fiber ports. It is suitable for wiring closet installations. It can also be equipped with up to 48 10 Gigabit Ethernet ports for use as the core switch.
The OmniSwitch 9600 chassis Note. OmniSwitch 9600 NI modules and OmniSwitch 7000 NI modules should not be mixed in the same Chassis.	The OmniSwitch 9600 chassis supports a high-performance switch fabric and provides 4 slots for Ethernet, Gigabit Ethernet, and/or 10 Gigabit Ethernet Network Interface (NI) modules. It provides one slot for primary Chassis Management Module (CMM)(there is no CMM redundancy in this chassis). The OmniSwitch 9600 supports a maximum of two power supplies and up to 96 10/100/1000 copper ports and/or 1000 Mbps fiber ports. It is suitable for wiring closet installations. It can also be equipped with up to 24 10 Gigabit Ethernet ports for use as the core switch.

System Requirements	Memory Requirements: • OmniSwitch 9000 Series Release 6.1.3.R01 requires 256 MB of SDRAM and 128MB of flash memory. This is the standard configuration shipped. Configuration files and the compressed software images—including web management software (WebView) images—are stored in the flash memory. Use the show hardware info command to determine your SDRAM and flash memory. Uboot, FPGA, MiniBoot, BootROM, and Upgrade (jfpga.upgrade_list: size; 1,211KB) Requirements: OmniSwitch 9000 Series • Uboot NI (size: 512KB): 6.1.1.167.R02 or later • Uboot CMM (size: 512KB): 6.1.1.167.R02 or later • MiniBoot. Uboot CMM (size: 843KB): 6.1.1.167.R02 or later • FPGA CMM: Major Revision: 2 Minor Revision: 25 (displays as 0x19) • software.lsm (size: 1KB) • POE Firmware: 5.01
Connections to the Chassis	Once your switch is properly installed, you should connect all network and management cables required for your network applications. Connections may include: • Serial cable to the console port • Ethernet cable to the Ethernet Management Port (EMP) on the CMM • Gigabit cables to all required XFPs or MiniGBICs • Ethernet cables to all required Ethernet Network Interface (ENI) ports
Chassis Management Module (CMM) **OS9600-CMM:** Chassis Management Module for the OS9600 & OS9700 **OS9700-CMM:** Chassis Management Module for the OS9600 & OS9700 **OS9800-CMM:** Chassis Management Module for the OS9800	The Chassis Management Module (CMM) is the management unit for OmniSwitch 9000 Series switches. It provides the main Switching Fabric & Management functionalities. **4.2.4, 4.2.6** In its role as the management unit, the CMM also provides key system services, including: • Console, USB, and Ethernet management port connections to the switch • Software and configuration management, including the Command Line Interface (CLI) • Web-based management (WebView) • SNMP management • Power Management & distribution • Temperature management • Switch diagnostics • Important availability features, including redundancy (when used in conjunction with another CMM), software rollback, temperature management, and power management • The CMM also contains the switch fabric unit for the OmniSwitch 9000. Data passing from one NI module to another passes through the CMM fabric. When two CMMs are installed, both fabrics are normally active. Note. The USB port on the front panel of the CMM is not supported in the 6.1.3r01 release. CMM Installation: On OmniSwitch 9000 Series switches, a minimum of one CMM is required for switch operations. On OmniSwitch 9700 & 9800, the second CMM provides redundancy. CMMs may be installed either in slot A or slot B in OS9700 & OS9800 switches. In non-redundant configurations, the CMM may be installed in either slot A or B. In redundant configurations, the CMM installed in slot A will be designated primary by default. NI modules cannot be installed in CMM slots A or B; likewise, CMMs cannot be installed in any NI slot position. Note. CMM redundancy is not supported on OmniSwitch 9600 switches because OS9600 contains only one CMM slot. *Note. OmniSwitch 9000 Series CMMs are colored **orange** to distinguish them from OmniSwitch 7700/7800 CMMs that are colored white. Do not install OmniSwitch 9000 Series and OmniSwitch 7700/7800 CMMs in the same chassis.*
CMM Redundancy **4.2.5**	CMM redundancy is an important resiliency feature. For CMM redundancy, two fully operational CMM modules must be installed in the chassis at all times. Note. CMM redundancy is not supported on OmniSwitch 9600 switches because OS9600 contains only one CMM slot. When two CMMs are running in the switch, one CMM has the primary role and the other has the secondary role at any given time. The primary CMM manages the current switch operations, while the secondary CMM provides backup (also referred to as “failover”). Note. By default, the CMM in slot A automatically assumes the primary role. If the primary CMM fails or goes offline for any reason, the secondary CMM is notified. The secondary CMM then automatically assumes the primary role.
Chassis Management Module (CMM) LEDs	The CMM provides a series of status LEDs on the module’s front panel. These LEDs offer basic status information for the following switch functions: • CMM hardware operation (OK1) • System software (OK2) • CMM processor status (CONTROL)

	• CMM fabric status (FABRIC) • Chassis ambient air temperature (TEMP) • Fan status (FAN) • Power Supply Unit (PSU) status (PSU) • Ethernet management port (LINK and ACT)
Component LEDs	Following a successful boot, the LEDs on all switch components, including power supplies, should display as follows: CMM OK1: Solid **Green** CMM OK2: Blinking **Green** CMM CONTROL: Solid **Green** CMM FABRIC: Solid **Green** CMM FAN: Solid **Green** CMM TEMP **Green** CMM PSU **Green** NI OK1: Solid **Green** NI OK2: Blinking **Green** Power Supply AC OK: Solid **Green** Power Supply DC OK: Solid **Green** Power Supply OVER TEMP: Off
NI Modules Installation	NI modules may be installed in any slot position from 1 through 16 in OS9800 switches, from 1 through 8 in OS9700 switches and 1 through 4 OS9600 switches.
OS9-GNI-C24 Module (10/100/1000BASE-T module)	• This module supports 24 x 10/100/1000BASE-T (10/100/1000Mbps) RJ45 ports. Each copper port is capable of auto-MDI/MDI-X sensing. The 10/100/1000BASE-T ports will operate in full/half duplex mode when the speed is 10/100Mbps. When operating in 1,000 Mbps only full duplex mode is supported.
OS9-GNI-C24 Module LEDs (10/100/1000BASE-T module)	**OK1:** Hardware Status. Displays solid green when powered on and the GNI has passed hardware diagnostic tests. Displays solid amber when powered on and the GNI has failed diagnostic testes. **OK2:** Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load the software. **PoE:** PoE Status. This LED will be off if PoE is not available on this module and will be solid green if PoE is enabled on this module. **Ethernet port LEDs:** Each Gigabit Ethernet port has two built-in corresponding LEDs. The top LED indicates 10/100Mbps link and activity status for the port while the bottom LED indicates 1 Gigabit link and activity status for the port. The appropriate LED displays solid green when a valid Ethernet cable connection exists and there is no PoE. Flashes green as data is transmitted or received on the port and there is no PoE. If PoE is present, the appropriate LED displays solid amber when a valid Ethernet cable connection exists. And flashes amber as data is transmitted or received on the port if PoE is present.
OS9-GNI-U24 Module (1000BASE-X module)	• This module supports 24 x 1000BASE-X (Gigabit Ethernet) SFP MSA fiber optics ports. The transceiver ports support full duplex mode only.
OS9-GNI-U24 Module LEDs	**OK1:** Hardware Status. Displays solid green when powered on and the GNI has passed hardware diagnostic tests. Displays solid amber when powered on and the GNI has failed diagnostic testes. **OK2:** Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load the software. **Gigabit Ethernet port LEDs:** Each fiber-based Gigabit Ethernet port has a corresponding LED. This LED indicates the link and activity status for each Gigabit Ethernet port. The LED displays green when a valid Gigabit Ethernet cable connection exists. Flashes green as data is transmitted or received on the port.
OS9-GNI-P24 Module (10/100/1000BASE-T module with PoE)	This module supports 24 x 10/100/1000BASE-T (10/100/1000Mbps) RJ45 ports with PoE (IEEE 802.3af). Each copper port is capable of auto-MDI/MDI-X sensing, twisted-pair Power over Ethernet (PoE) ports, individually configurable as 10Base-T, 100Base-TX, or 1000Base-T. The 10/100/1000BASE-T ports will operate in full/half duplex mode when the speed is 10/100Mbps. When operating in 1,000 Mbps only full duplex mode is supported.
OS9-GNI-P24 Module LEDs (10/100/1000BASE-T module with PoE)	**OK1:** Hardware Status. Displays solid green when powered on and the GNI has passed hardware diagnostic tests. Displays solid amber when powered on and the GNI has failed diagnostic testes. **OK2:** Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load the software. **PoE:** PoE Status. This LED will be off if PoE is not available on this module and will be solid green if PoE is enabled on this module. **Ethernet port LEDs:** Each Gigabit Ethernet port has two built-in corresponding LEDs. The top LED indicates 10/100Mbps link and activity status for the port while the bottom LED indicates 1 Gigabit link and activity status for the port. The appropriate LED displays solid green when a valid Ethernet cable connection exists and there is no PoE. Flashes green as data is transmitted or received on the port and there is no PoE. If PoE is present, the appropriate LED displays solid amber when a valid Ethernet cable connection exists. And flashes amber as data is transmitted or received on the port if PoE is present.

OS9-XNI-U2 Module (2 x 10-Gigabit Ethernet XFP NI module)	OS9-XNI-U2: 2 x XFP 10-GigEth ports Each 10 GigEth port supports industry standard XFP based 10GigE SMF 10GBASE-LR, MMF 10GBASE-SR, SMF 10GBASE-ER, and SMF 10BASE-ZR optical transceivers. The applicable models provide 2 XFP slots. These slots support the following XFP types: ▪ XFP-10G-ER40—10GBASE-ER Single mode fiber, supports distances up to 40km; uses LC connectors. ▪ XFP-10G-LR—10GBASE-LR Single mode fiber, supports distances up to 10km; uses LC connectors. ▪ XFP-10G-SR—10GBASE-SR Multimode fiber, supports distances up to 300m; uses LC connectors. ▪ XFP-10G-ZR80—10GBASE-ZR Single mode fiber, supports distances up to 80km; uses LC connectors. The two-port XFP 10 Gigabit slots can mix and match different 10-Gigabit XFP transceiver types. *Note. Compatibility with the OmniSwitch 6800 & OmniSwitch 6850 10-Gigabit Ethernet is supported.*
OS9-XNI-U2 LEDs Module (2 x 10-Gigabit Ethernet XFP NI module)	**OK1:** Hardware Status. Displays solid green when powered on and the GNI has passed hardware diagnostic tests. Displays solid amber when powered on and the GNI has failed diagnostic testes. **OK2:** Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load the software. LINK/ACT LED Each 10-Gigabit port has a single LED for monitoring XFP link status and activity. The LED displays solid green when the port is up; the LED blinks green when the port is transmitting or receiving packets in a link up state. The LED is off when no link is detected.
OS9-XNI-U6 Module (6 x 10-Gigabit Ethernet XFP NI module)	OS9-XNI-U6: 6 x XFP 10-GigEth ports Each 10 GigEth port supports industry standard XFP based 10GigE SMF 10GBASE-LR, MMF 10GBASE-SR, SMF 10GBASE-ER, and SMF 10BASE-ZR optical transceivers. The applicable models provide 2 XFP slots. These slots support the following XFP types: ▪ XFP-10G-ER40—10GBASE-ER Single mode fiber, supports distances up to 40km; uses LC connectors. ▪ XFP-10G-LR—10GBASE-LR Single mode fiber, supports distances up to 10km; uses LC connectors. ▪ XFP-10G-SR—10GBASE-SR Multimode fiber, supports distances up to 300m; uses LC connectors. ▪ XFP-10G-ZR80—10GBASE-ZR Single mode fiber, supports distances up to 80km; uses LC connectors. The six-port XFP 10 Gigabit slots can mix and match different 10-Gigabit XFP transceiver types. *Note. Compatibility with the OmniSwitch 6800 & OmniSwitch 6850 10-Gigabit Ethernet is supported*
OS9-XNI-U6 Module LEDs (6 x 10-Gigabit Ethernet XFP NI module)	**OK1:** Hardware Status. Displays solid green when powered on and the GNI have passed hardware diagnostic tests. Displays solid amber when powered on and the GNI have failed diagnostic testes. **OK2:** Software Status. Blinks green when the GNI is operational and has successfully loaded software. Displays solid amber when powered on and the GNI has failed to load the software. LINK/ACT LED Each 10-Gigabit port has a single LED for monitoring XFP link status and activity. The LED displays solid green when the port is up; the LED blinks green when the port is transmitting or receiving packets in a link up state. The LED is off when no link is detected.
Transceivers	
10 Gigabit Ethernet Transceivers (XFP MSA)	
10-Gigabit Ethernet XFP	XFP 10-GigEth ports Each 10 GigEth port supports industry standard XFP based 10GigE SMF 10GBASE-LR, MMF 10GBASE-SR, SMF 10GBASE-ER, and SMF 10BASE-ZR optical transceivers. The applicable models provide 2 XFP slots. These slots support the following XFP types: ▪ XFP-10G-ER40—10GBASE-ER Single mode fiber, supports distances up to 40km; uses LC connectors. ▪ XFP-10G-LR—10GBASE-LR Single mode fiber, supports distances up to 10km; uses LC connectors. ▪ XFP-10G-SR—10GBASE-SR Multimode fiber, supports distances up to 300m; uses LC connectors. ▪ XFP-10G-ZR80—10GBASE-ZR Single mode fiber, supports distances up to 80km; uses LC connectors. The XFP 10 Gigabit slots can mix and match different 10-Gigabit XFP transceiver types. *Note. Compatibility with the OmniSwitch 6800 & OmniSwitch 6850 10-Gigabit Ethernet is supported.*
XFP-10G-ER40	10 Gigabit Ethernet optical transceiver (XFP MSA): Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 40km on 9/125µm SMF.
XFP-10G-LR	10 Gigabit Ethernet optical transceiver (XFP MSA): Supports single mode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 10km on 9/125µm SMF. [Formerly known as 10G-XFP-LR]
XFP-10G-SR	10 Gigabit Ethernet optical transceiver (XFP MSA):

	Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Typical reach of 300m on 50/125µm MMF. [Formerly known as 10G-XFP-SR]
XFP-10G-ZR80	10 Gigabit Ethernet optical transceiver (XFP MSA): Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 80km on 9/125µm SMF.
Gigabit Ethernet Transceivers (SFP MSA)	
SFP-GIG-47CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ gray latch. Supports single mode fiber over 1470 nm wavelength (nominal) with an LC connector. Typical reach of 62Km on 9/125µm SMF.
SFP-GIG-49CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ violet latch. Supports single mode fiber over 1490 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125µm SMF.
SFP-GIG-51CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ blue latch. Supports single mode fiber over 1510 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125µm SMF.
SFP-GIG-53CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ green latch. Supports single mode fiber over 1530 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125µm SMF.
SFP-GIG-55CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ yellow latch. Supports single mode fiber over 1550 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125µm SMF.
SFP-GIG-57CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ orange latch. Supports single mode fiber over 1570 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125µm SMF.
SFP-GIG-59CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ red latch. Supports single mode fiber over 1590 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125µm SMF.
SFP-GIG-61CWD60	CWDM Gigabit Ethernet optical transceiver (SFP MSA) w/ red latch. Supports single mode fiber over 1610 nm wavelength (nominal) with an LC connector. Typical reach of 62 Km on 9/125µm SMF.
SFP-GIG-EXTND	Extended 1000Base-SX Gigabit Ethernet optical transceiver (SFP MSA): Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Reach of up to 2 km on 62.5/125µm MMF and 50/125µm MMF. Requires SFP-GIG-EXTND or GBIC-GIG-EXTND at the remote termination. [Formerly known as GE-EXTND-SFP]
SFP-GIG-LH40	1000Base-LH40 Gigabit Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1310 nm wavelength (nominal) with an LC connector. Typical reach of 40Km on 9/125µm SMF.
SFP-GIG-LH70	1000Base-LH70 Gigabit Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1550nm wavelength (nominal) with an LC connector. Typical reach of 70 Km on 9/125µm SMF. [Formerly known as MINIGBIC-LH-70]
SFP-GIG-LX	1000Base-LX Gigabit Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 10 Km on 9/125µm SMF. Typical reach of 550m on 50/125 & 62.5/125µm MMF. [Formerly known as MINIGBIC-LX]
SFP-GIG-SX	1000Base-SX Gigabit Ethernet optical transceiver (SFP MSA): Supports multimode fiber over 850nm wavelength (nominal) with an LC connector. Typical reach of 300m on 62.5/125µm MMF or 550m on 50/125µm MMF. [Formerly known as MINIGBIC-SX]
SFP-GIG-T	1000Base-T Gigabit Ethernet Transceiver (SFP MSA) - Supports category 5, 5E, and 6 copper cabling up to 100m. SFP only works in 1000 Mbps speed and full-duplex mode.
Dual Speed Ethernet Transceivers (SFP MSA)	
SFP-DUAL-MM	Dual Speed 100Base-FX or 1000Base-X Ethernet optical transceiver (SFP MSA). Supports multimode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 550m at Gigabit speed and 2km at 100Mbit speed. Notes: - At 100Mbit speed, this SFP can interoperate with SFP-100-LC-MM or similar transceiver on the other end - At Gigabit speed, this SFP cannot interoperate with SFP-GIG-SX or similar transceiver on the other end-running over 850nm wavelength. It can interoperate with 1000-BASE-LX over MMF or another SFP-DUAL-MM. - SFP supported on OS9-GNI-U24 Gig. Ethernet Module and OS6850-U24X SFP ports (non combo)
SFP-DUAL-SM10	Dual Speed 100Base-FX or 1000Base-X Ethernet optical transceiver (SFP MSA). Supports single mode fiber over 1310nm wavelength (nominal) with an LC connector. Typical reach of 10km at Gigabit speed and 100Mbit speed. Notes: - At 100Mbit speed, this SFP can interoperate with SFP-100-LC-SM15 or similar transceiver, - At Gigabit speed, this SFP can interoperate with SFP-GIG-LX or similar transceiver. - SFP supported on OS9-GNI-U24 Gig. Ethernet Module and OS6850-U24X SFP ports (non combo)

Supported Configuration Matrix for New Ethernet Transceivers in Release 6.1.3r01

SFP	OS6800/OS6850 Combo Ports	OS6800-U24 Non-Combo Ports	OS6850-U24X	OS9-GNI-U24
SFP-GIG-T - 1000Base-T Gigabit Ethernet Transceiver (SFP MSA).	**Supported**	**Supported**	**Supported**	**Supported**

SFP-DUAL-MM - Dual Speed 100Base-FX or 1000Base-X Ethernet optical transceiver.	**Un-supported**	**Un-supported**	**Supported**	**Supported**
SFP-DUAL-SM10 - Dual Speed 100Base-FX or 1000Base-X Ethernet optical transceiver (SFP MSA)	**Un-supported**	**Un-supported**	**Supported**	**Supported**
SFP-100-BX20LT - 100Base-BX SFP bi-directional transceiver.	**Un-supported**	**Un-supported**	**Supported**	**Un-supported**
SFP-100-BX20NU - 100Base-BX SFP bi-directional transceiver.	**Un-supported**	**Un-supported**	**Supported**	**Un-supported**
SFP-100-LC-MM - 100Base-FX SFP transceiver.	**Un-supported**	**Un-supported**	**Supported**	**Un-supported**
SFP-100-LC-SM15 - 100Base-FX SFP transceiver.	**Un-supported**	**Un-supported**	**Supported**	**Un-supported**
SFP-100-LC-SM40 - 100Base-FX SFP transceiver.	**Un-supported**	**Un-supported**	**Supported**	**Un-supported**
Hardware Architecture				
Network Interface	All modules are hot swappable and can be used in any available NIs slot **4.2.7** 24--port 1000BaseX (SFP) 24-port 10/100/1000 BASE-T (RJ45) 24-port PoE 10/100/1000BaseT (RJ45) 2-port 10GBaseX (XFP) 6-port 10GBaseX (XFP)			
MAC Address Table	In synchronized mode (default), up to 16K MAC Addresses is supported per system In Distributed mode, up to 64 K MAC Addresses is supported per system (no more than 16K per NI). 1K (authenticated / mobile users) per system Latency: <10µsec			
Learned Port Security – What is the maximum number of MAC addresses a port can learn?	For the OmniSwitch 9000 family, the learned port security feature of the Alcatel.Lucent Operating System allows up to 100 MAC addresses per port to be learned and acted upon.			
IP Address Table Routes	40K routing table 12K forwarding LPM entries, 8K hosts entries per module Latency: <10µsec			
Manufacturing MAC Address Assignments	Manufacturing will allocate a group of MAC addresses from the Alcatel.Lucent range of IEEE registered MAC addresses for each OS9000 backplane built. The block is a set of 32 consecutive MAC addresses for each newly built OS9000 system. These MAC addresses are stored on the backplane EEPROMs. The Ethernet Management Port (EMP) uses a MAC address taken from the backplane PROMs.			
CPU	Motorola/Free-scale MPC8540/MPC8541 Power PC.			
BUS	FBUS+; between the Fabric Board and the NIs. Each FBUS+ allowing 12Gbps full duplex			
Memory	256MB (DRAM) per CMM			
Flash	128MB per CMM Note: the OS9000 will use USB (future release) to provide expandable flash memory.			
Switching MAC ASIC	BCM5650			
Main Switching Fabric ASIC	On the CMM: BCM5675 and on each NI: BCM56504s and BCM5650s The BCM5675 single chip fabric provides eight queues for each egress port on that chip. Because, the BCM5675 has eight fabric ports, one for each of the eight NI/line card slots (OS9700 chassis), the fabric has 8x8=64 unicast queues. Each of the eight egress queues represents a different class of service (COS) or priority. Various strict priority and round-robin bandwidth control options can be configured to provide for DiffServ PHBs (per hop behavior).			
10-Gigabit Ethernet Interface	10-Gigabit Ethernet XAUI interface			
PHY	On the OS9-GNI-C24 & OS9-GNI-P24 NIs: BCM5464R			
Packet Buffer Size per NI	2MB			
RS-232 Console Port	RJ-45 connector (please refer to the section “Pin-Outs” for further details.			
EEPROM	Up to 2K			
Front Panel LED	Each NI Module and Power Supply supports appropriate Status type LEDs.			
Temperature Sensor	Temperature Sensor is supported			
Power Supply	AC-to-DC and DC-to-DC N+1 redundant Power Supplies are supported. **4.2.8** OS9600 chassis type supports up to 2 such power supplies. OS9700 chassis type supports up to 3 such power supplies. OS9800 chassis type supports up to 4 such power supplies.			
Ethernet Specifications				
Connectors/ Cabling	• Management: 1 RJ-45 console interface configured as DCE/DTE for operation, diagnostics, status, and configuration information. Ship kit includes RJ-45 to DB-9 connector adaptor • AC power connector			

Connector type	10/1000/1000BASE-T copper ports without PoE: RJ-45 10/100/1000BASE-T copper ports with PoE: RJ-45 Dual Speed 100BASE-FX or 1000BASE-X SFP ports: LC w/ Removable/Pluggable trans. SFP-MSA CWDM SFP ports: LC with Removable/Pluggable transceiver – SFP-MSA 1000BASE-X SFP ports: LC with Removable/Pluggable transceiver – SFP-MSA 1000BASE-T SFP ports: LC with Removable/Pluggable transceiver – SFP-MSA 10GBASE-X XFP ports: LC with Removable/Pluggable XFP-MSA transceiver Please refer to "Supported Configuration Matrix for New Ethernet Transceivers in Release 6.1.3r01"
Connectivity	All modules are hot swappable and can be used in any available NI slot. 10/1000/1000BASE-T copper ports: RJ-45 10/100/1000BASE-T copper ports with PoE: RJ-45 100BASE-FX or 1000BASE-X SFP ports: LC with Removable/Pluggable transceiver – SFP-MSA CWDM SFP ports: LC with Removable/Pluggable transceiver – SFP-MSA 1000BASE-X SFP ports: LC with Removable/Pluggable transceiver – SFP-MSA 1000BASE-T SFP ports: LC with Removable/Pluggable transceiver – SFP-MSA 10GBASE-X XFP ports: LC with Removable/Pluggable XFP-MSA transceiver
Connections supported	CWDM Gigabit Ethernet Dual Speed 100BASE-FX or 1000BASE-X SFP ports 10BASE-T hub or device; 100BASE-TX hub or device; 1000BASE-T hub or device 1000BASE-X hub or device, and 10GBASE-X hub or device
Cable supported	CWDM Gigabit Ethernet, Dual Speed 100BASE-FX or 1000BASE-X, 1000BASE-X, and 10GBASE-X: Optical Fiber 10BASE-T: unshielded twisted-pair (UTP) 100BASE-TX: unshielded twisted-pair (UTP), Category 5, EIA/TIA 568 or shielded twisted-pair (STP), Category 5, 100 ohm 1000BASE-T: unshielded twisted-pair (UTP), Category 5/5e, EIA/TIA 568 or shielded twisted-pair (STP), Category 5, 100 ohm *Note:* Category 6 cabling is also supported on the 10/100/1000BASE-T connections.
Maximum cable distance	On 10/100/1000Mbps triple speed copper ports: 10Mbps speed: 100 meters on copper 100Mbps speed: 100 meters on copper 1000Mbps speed: 100 meters on copper On GigE. Fiber ports: SFP-GIG-xCWD60: up to 60km on 9/125µm SMF SFP-GIG-LH40: up to 40km on 9/125µm SMF SFP-GIG-LH70: up to 70km on 9/125µm SMF SFP-GIG-LX: 10km on 9/125µm SMF & typical reach of 550m on 50/125 & 62.5/125µm MMF. SFP-GIG-SX: up to 550m on 50/125 & 62.5/125µm MMF SFP-GIG-T: up to 100m on copper Dual Speed Fiber ports: SFP-DUAL-MM: 550m @Gigabit speeds and 2km @100Mbps speeds - MMF SFP-DUAL-SM10: 10km @Gigabit speeds and @100Mbps speeds - SMF On 10GigE. Fiber ports: • XFP-10G-SR: up to 300 m (high modal bandwidth fiber is required to reach 300 meters) • XFP-10G-LR: up to 10 km • XFP-10G-ER40: up to 40km • XFP-10G-ZR80: up to 80km
IEEE Standards Supported	IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD)
Data rates	• Dual Speed 100Mbps or Gigabit speeds (Dual Speed 100BASE-FX or 1000BASE-X) • 10/100/1000Mbps triple speed o 10Mbps o 100Mbps o 1000Mbps (Gigabit Ethernet) • Gigabit Ethernet • 10000Mbps (10-Gigabit Ethernet)
Ports Supported **4.2.9**	• Dual Speed 100Mbps or Gigabit speeds (Dual Speed 100BASE-FX or 1000BASE-X) • Triple Speed ports is supported and includes: o Ethernet (10 Mbps) o Fast Ethernet (100 Mbps) o 1000Mbps Ethernet (Gigabit Ethernet) • Gigabit Ethernet • 10-Gigabit Ethernet
Switching/Routing Support	Layer 2 Switching/Layer 3 Routing
Backbone Support	100Mbps/1000Mbps, 10/100/1000Mbps, Gigabit Ethernet ports, and 10-Gigabit Ethernet ports
Port Mirroring Support	100Mbps/1000Mbps, 10/100/1000Mbps, Gigabit Ethernet ports, and 10-Gigabit Ethernet ports
802.1Q Hardware Tagging	100/1000Mbps, 10/100/1000Mbps, Gigabit Ethernet ports, and 10-Gigabit Ethernet ports
Maximum Transfer Unit -- MTU	The ASIC does not include the notion of an MTU that applies to an IP interface. Instead, it uses the

	physical long-frame-size of the egress port as the MTU. When the ASIC attempts to forward a packet, it tests the size of the packet against the physical long-frame-size of the egress port, if the packet is too large, it forwards the packet to the CPU for fragmentation (or ICMP processing in the case of a packet with Do not Fragment set). • 10/100 ports are set with a long-frame-size of 1553 bytes. • GigE/10GigE ports are set with a long-frame-size of 9216 bytes (jumbo frames). **4.1.7** Packets larger than the long-frame-size are dropped at ingress. The above (& default) values are the maximum configurable values. Packets that are forwarded from a 10/100 to a 10/100 port cannot ever be reported as too big via ICMP because anything larger than 1553 would not be accepted. The same holds true for packets forwarded between two GigE/10GigE ports and from a 10/100 port to a GigE/10GigE. Layer-2 Ethernet Frame Size: Untagged: 1,518 Bytes without IEEE 802.1Q tags Tagged: 1,522 Bytes with IEEE 802.1Q tags Long Frame Size (enabled by default): 1553 Bytes (IEEE 8021.Q tagged or untagged) Frame Type: Type2, LLC, SNAP, RAW 802.3 The maximum frame size on the Gigabit Ethernet interfaces range from 1,518 to 9,216 Bytes Jumbo frames up to 9K Bytes (9,216 Bytes) are supported on ALL Module types including the OS9-XNI-U2 & OS9-XNI-U6 (10-Gigabit Ethernet Modules). Untagged (without IEEE 802.1Q tags) Ethernet Packets: 1,518 Bytes Tagged (with IEEE 802.1Q tags) Ethernet Packets: 1,522 Bytes
Inter-Frame Gap	12 Bytes (by default) Inter-frame gap is a measure of the minimum idle time between the end of one frame transmission and the beginning of another. By default, the inter-frame gap is 12 bytes. Through the use of "interfaces ifg" Command, the inter-frame gap value (in bytes) on a specific port, a range of ports, or all ports on a switch (slot) can be configured. Values for this command range from 9 to 12 bytes. **Note.** This command is only valid on Gigabit ports.
Interface Alias (Port Alias)	Supported (none configured by default): Through the use of this feature an alias (i.e., description) for a single port can be configured. (You cannot configure an entire switch or a range of ports.) The text description can be up to 40 characters long.
Peak Flood Rate Configuration	By default: 4 Mbps (10 Ethernet) 49 Mbps (100 Fast Ethernet) 496 Mbps (1 Gigabit Ethernet) 997 Mbps (10 Gigabit Ethernet) Through the use of this feature, the peak ingress flood rate value on a specific port, a range of ports, or all ports on a switch (slot) in megabits per second can be configured. **Note. The user** can configure a flood rate equal to the line rate, but it is not recommended. Alcatel.Lucent recommends that you always configure the flood rate to be less than the line speed.
Flow Control	The **flow** command can be used to enable (the default) or disable flow control on a specific port, a range of ports, or all ports on an entire switch (slot). When the buffers on a receiving device are full, flow control transmits pause frames to the remote link partner to delay transmission. The local port can delay transmission of data if the remote link partner transmits a pause frame. By default, the flow control wait time is 0 microseconds. IEEE 802.3x (programmable threshold) flow control. – Enabled by default (Note: the switch supports and honors the incoming IEEE 802.3x pause frames, but it does not generate outgoing IEEE 802.3x pause frames)
Trap Port Link Messages	Supported (disabled by default) This feature can be enabled or disabled (the default) on a specific port, a range of ports, or all ports on a switch (slot). When enabled, a trap message will be displayed on a Network Management Station (NMS) whenever the port state has changed.
Per port rate limiting Per-port L2/L3 multicast & broadcast flood limit is supported.	Per-port multicast / broadcast / flood limit is supported. The ASIC provides a per port configuration on the incoming and/or outgoing port basis that allows broadcast and/or multicast storm control. The CPU can program a threshold value per port that indicates the number of broadcast and/or multicast packets/bytes that are allowed in a given time interval.
Re-settable Statistics Counters	Supported
Duplex Mode support	The duplex mode feature is supported on a specific port, a range of ports, or all ports on a switch (slot). It can be set to **full** (full duplex mode, which is the default on fiber ports), **half** (half duplex mode), and **auto** (auto-negotiation, which is the default on copper ports). The **Auto** option causes the switch to advertise all available duplex modes (half/full/both) for the port during auto-negotiation. In full duplex mode, the interface transmits and receives data simultaneously. In half duplex mode, the interface can only transmit or receive data at a given time.
Auto-negotiation	Auto-negotiation is supported (enabled by default). It can be enabled or disabled on a single port, a range of ports, or an entire slot.
Crossover	Crossover can be configured on a single port, a range of ports, or an entire slot. If auto negotiation is disabled, auto MDIX, flow control, auto speed, and auto duplex are not accepted.

	Setting the crossover configuration to **auto** will configure the interface or interfaces to automatically detect crossover settings. Setting crossover configuration to **mdix** will configure the interface or interfaces for MDIX (Media Dependent Interface with Crossover), which is the standard for hubs and switches. Setting crossover to **mdi** will configure the interface or interfaces for MDI (Media Dependent Interface), which is the standard for end stations. And setting the crossover configuration to **disable** will disable crossover configuration on an interface or interfaces.
Verifying Ethernet Port Configurations	To display information about Ethernet port configuration settings, use the **show** commands. These commands can be quite useful in troubleshooting and resolving potential configuration issues or problems on your switch. For more information about the resulting displays from these commands, see the *OmniSwitch CLI Reference Guide*.
Performance	
Performance	
Principle of operation for Fabric Load Sharing	• Traffic intra-module to be processed & forwarded locally • Traffic inter-module to be forwarded through the Virtual Switching Fabric • Each Chassis Management Module (CMM) provides a 12Gbps Full Duplex (24Gbps aggregated) connection between each NI module. For OS9600, each NI module leverages 2 connection to the CMM (versus 1 connection per CMM in OS9700/OS9800) • OmniSwitch 9700 & 9800 full switching capacity is reached with dual CMMs • OmniSwitch 9600 full switching capacity is reached with a single CMM
Principle of operation for Distributed Processing	• Each NI module provides a high performance CPU o CMM's CPU is responsible for management & overall coordination o NI Module's CPU is responsible for most operations • Management bus is a dedicated Gigabit Ethernet Full Duplex bus o Each NI module's CPU supports a direct connection with each CMM's CPU
Switching Scheme supported	Locally within each NI Module or through the Switch Fabric on the CMM
Forwarding Capabilities Performances • Centrally through the Switch Fabric on the CMM OR • Locally within the same NI Module	Each NI Module is capable of switching (forwarding) traffic **centrally or locally**: **OS9-GNI-C24:** 24Gbps Full Duplex (aggregated bandwidth of 48Gbps) Switching (Forwarding) Bandwidth Up to 35.7Mpps switching (forwarding) Throughput, allowing wire-speeds operation. **OS9-GNI-U24:** 24Gbps Full Duplex (aggregated bandwidth of 48Gbps) Switching (Forwarding) Bandwidth Up to 35.7Mpps switching (forwarding) Throughput, allowing wire-speeds operation. **OS9-GNI-P24:** 24Gbps Full Duplex (aggregated bandwidth of 48Gbps) Switching (Forwarding) Bandwidth Up to 35.7Mpps switching (forwarding) Throughput, allowing wire-speeds operation. **OS9-XNI-U2:** 24Gbps Full Duplex (aggregated bandwidth of 48Gbps) Switching (Forwarding) Bandwidth Up to 29.7Mpps switching (forwarding) Throughput, allowing wire-speeds operation. **OS9-XNI-U6:** The OS9-XNI-U6 module (6-port 10GBASE-X) supports 6 x 10-GigE ports with an oversubscription ratio of: 2.5:1 (6 x 10GigE / 24 GigE full duplex bandwidth per slot) 24Gbps Full Duplex (aggregated bandwidth of 48Gbps) Switching (Forwarding) Bandwidth Up to 35.7Mpps switching (forwarding) Throughput Note: The OS9-XNI-U6 module supports local switching, allowing wire-speeds operation under specific conditions: Local switching between ports: 1, 2 and 3 is supported at wire- speed (44.6 Mpps) Local switching between ports: 4, 5 and 6 is supported at wire- speed (44.6 Mpps)
Backplane Architecture & Backplane Capacity	The OmniSwitch 9000 Series support a Passive Backplane. ❑ **OS9600: capable of 960Gbps** ❑ **OS9700: capable of 960Gbps** **Calculation Method for both OS9600 & OS9700 Backplane Arch. & Backplane Capacity:** Backplane capacity of 960Gbps -Today, we use 4 lanes going in each direction, to each fabric, out of 8 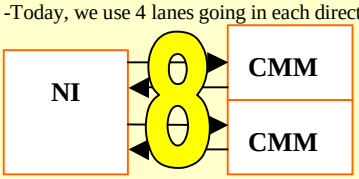 -Each lane is running at 3.75GHz - 3.75GHz x 4 = 15GHz (15Gbps) - 8B/10B encoding => 12Gbps efficient -If we are running ALL lanes at that speed then: -Each lane, is going to each CMM, in each direction, on 8 slot @ 3.75Gbps - 8 x 2 x 2 x 8 x 3.75 = 960Gbps

	❑ **OS9800 Backplane Arch. & Backplane Capacity: capable of 1.92Tbps** Backplane capacity of 1.92Tbps -Today, we use 4 lanes going in each direction, to each fabric, out of 8 -Each lane is running at 3.75GHz - 3.75GHz x 4 = 15GHz (15Gbps) - 8B/10B encoding => 12Gbps efficient -If we are running ALL lanes at that speed then: -Each lane, is going to each CMM, in each direction, on 16 slot @ 3.75Gbps - 8 x 2 x 2 x 16 x 3.75 = 1.92Tbps
Architecture **Note: All NI Modules support local switching.**	The OS9700 & OS9800 system uses two fabric cards in load sharing mode to provide full system capacity. The OS9600 system uses one fabric card to provide full system capacity. Full wire rate means that every port is sending and receiving packets continuously at the maximum Gigabit Ethernet or Ten Gigabit Ethernet rates. OS9700 & OS9800 provides full wire rate to all user ports by distributing traffic evenly between the two fabric cards based on the L2, L3 and L4 addresses. OS9600 provides full wire rate to all user ports by distributing traffic evenly on one fabric card based on the L2, L3 and L4 addresses. **The OS9-GNI-C24 module (24-port 10/100/1000BASE-T RJ45 module) supports up to 24 GigE ports at wire-speeds.** **The OS9-GNI-P24 module (24-port 10/100/1000BASE-T RJ45 module with PoE) supports up to 24 GigE ports at wire-speeds.** **The OS9-GNI-U24 module (24-port 1000BASE-X) supports up to 24 GigE ports at wire-speeds.** **The OS9-XNI-U2 module (2-port 10GBASE-X) supports 2 x 10-GigE ports at wire-speed.** **The OS9-XNI-U6 module (6-port 10GBASE-X) supports 6 x 10-GigE ports with an oversubscription ratio of: 2.5:1 (6 x 10GigE / 24 GigE full duplex bandwidth per slot)** **Note:** **The OS9-XNI-U6 module supports local switching, allowing wire-speeds operation under the following specific conditions:** **Local switching between ports: 1, 2 and 3 is supported at wire-speed** **Local switching between ports: 4, 5 and 6 is supported at wire-speed.** **The OmniSwitch 9600 supports up to 96Gigabit Ethernet ports without PoE at wire-speed.** **The OmniSwitch 9600 supports up to 96 Gigabit Ethernet ports with PoE at wire-speed.** **The OmniSwitch 9600 supports up to 8 x 10-Gigabit Ethernet ports at wire-speed.** **The OmniSwitch 9600 supports up to 24 x 10-Gigabit Ethernet ports at 2.5:1 Oversubscription with QoS implementation** **The OmniSwitch 9700 supports up to 192 Gigabit Ethernet ports without PoE at wire-speed.** **The OmniSwitch 9700 supports up to 192 Gigabit Ethernet ports with PoE at wire-speed.** **The OmniSwitch 9700 supports up to 16 x 10-Gigabit Ethernet ports at wire-speed.** **The OmniSwitch 9700 support s up to 48 x 10-Gigabit Ethernet ports at 2.5:1 Oversubscription with QoS implementation** **The OmniSwitch 9800 supports up to 384 Gigabit Ethernet ports without PoE at wire-speed.** **The OmniSwitch 9800 supports up to 384 Gigabit Ethernet ports with PoE at wire-speed.** **The OmniSwitch 9800 supports up to 32 x 10-Gigabit Ethernet ports at wire-speed.** **The OmniSwitch 9800 support s up to 96 x 10-Gigabit Ethernet ports at 2.5:1 Oversubscription with QoS implementation**
Bandwidth per switch slot	**OmniSwitch 9700 & 9800 chassis:** 12Gbps Full Duplex or aggregated bandwidth of 24Gbps (NI ↔ Switch Fabric) with single CMM 24Gbps Full Duplex or aggregated bandwidth of 48Gbps (NI ↔ Switch Fabric) with dual CMM **OmniSwitch 9600 chassis:** 2 x 12Gbps Full Duplex or aggregated bandwidth of 48Gbps (NI ↔ Switch Fabric) with single CMM Note that the design of the OS9600 chassis will route two fabric interfaces (2 x 12Gbps FD) to each of the four slots (while an OS9700 will route a single Fabric interface to each of its 8 slots).
Throughput Performance Or Forwarding Rate Per Switch Or Maximum aggregated throughput	The OmniSwitch 9600/9700/9800 can forward packets on all ports simultaneously at full wire rate, even when the packets are minimum length. The forwarding rate is therefore: ❑ **OS9600: L2/L3/L4 Forwarding Rate of 142.85Mpps with 64Byte packets** ❑ **OS9700: L2/L3/L4 Forwarding Rate of 285.7Mpps with 64Byte packets** 4.2.2, 4.2.3 ❑ **OS9800: L2/L3/L4 Forwarding Rate of 571.4Mpps with 64Byte packets** **Calculation Method:** **The OmniSwitch 9600 fully loaded supports up to 96 Gigabit Eth ports at wire-speeds:** **96 * 1,488,095.23 pps = 142,857,142.1pps (142.85Mpps)**

	The OmniSwitch 9700 fully loaded supports up to 192 Gigabit Eth ports at wire-speeds: 192 * 1,488,095.23 pps = 285,714,284.2pps (285.7Mpps) **The OmniSwitch 9800 fully loaded supports up to 384 Gigabit Eth ports at wire-speeds: 384 * 1,488,095.23 pps = 571,428,568.3pps (571.4Mpps)**
Throughput Performance Or Forwarding Rate Per Module	**The OS9-GNI-C24 module (24-port 10/100/1000BASE-T RJ45 module) supports up to 24 GigE ports at wire-speeds: 24 * 1,488,095.23 pps = 35,714,285.52pps (35.7Mpps)** **The OS9-GNI-P24 module (24-port 10/100/1000BASE-T RJ45 module with PoE) supports up to 24 GigE ports at wire-speeds: 24 * 1,488,095.23 pps = 35,714,285.52pps (35.7Mpps)** **The OS9-GNI-U24 module (24-port 1000BASE-X SFP module) supports up to 24 GigE ports at wire-speeds: 24 * 1,488,095.23 pps = 35,714,285.52pps (35.7Mpps)** **The OS9-XNI-U2 module (2-port 10GBASE-X XFP module) supports 2 x 10-GigE ports at wire-speed: 2 * 14,880,952.3 pps = 29,761,904.6pps (29.7Mpps)** **The OS9-XNI-U6 module (6-port 10GBASE-X) supports 6 x 10-GigE ports with an oversubscription ratio of: 2.5:1. But, it can sustain 24 GigE ports at wire-speed (Bandwidth per switch slot: 24Gbps Full Duplex). Therefore, for throughput calculations: 24 * 1,488,095.23 pps = 35,714,285.52pps (35.7Mpps)**
Layer-2 & Layer-3 Forwarding Rate Per port	Wire-speed on 10Mbps port→ 14,880 pps with 64 Byte packets Wire-speed on 100Mbps port→ 148,809 pps with 64 Byte packets Wire-speed on Gigabit Ethernet port→ 1,488,095 pps with 64 Byte packets Wire-speed on 10-Gigabit Ethernet port→ 14,880,952 pps with 64 Byte packets
Latency	<10µsec with 64Byte packets
System	
Boot time	For all three platforms; OmniSwitch 9600/9700/9800 Cold boot time in a fully loaded configuration: approximately 60 sec. Warm re-boot time in a fully loaded configuration: approximately 60 sec.
Uboot Bootup Process	The Uboot code is responsible for loading the system kernel. Uboot is replacing the legacy MiniBoot/BootROM. The Uboot resides in the NVRAM on each CMM and NI. In the event where the Uboot process fails, most likely because the system can not load the Jos.img (i.e. image file getting corrupted), there is a backup mechanism to revert to MiniBoot. In the failed case, the switch will stop at the MiniBoot prompt and it will allow the user to do the following - reconfigure the EMP port - Ftp/zmodem the images in order to recover the corrupted files. Once the image files are recovered, the user can reload the box through the regular Uboot process. Please refer to the upgrade instructions guide for instructions on updating u-boot version.
Fabric load balance	To support wire speed capability on a fully loaded OS9800 & OS9700 chassis, you must have 2 CMM/Fabric. Each NI has one 12Gbps FD link to each fabric. To support wire speed from NI to Fabric the traffic is load balanced across the two 12Gbps FD links using the same algorithm as LinkAgg. Unlike link aggregation, flooding and IP multicast is always load balanced across the 2 fabrics using the same rule: mac-sa / mac-da for non ip packets or ip-sa / ip-da for ip packets
Management fail-over	The Fail-over time (Primary CMM to Secondary CMM) is in sub-second. In actuality the failover time is estimated to be around 70 to 300msec. Trap is sent (to the management station for the failure of the primary management) and log event is logged upon primary management failure and after the redundant management unit takes over. The CMM module contains hardware and software elements to provide management functions for the OS9000 system. Each CMM consists of two sub-modules · Processor module (CPM) · Fabric module (CFM) The OS9700 & OS9800 will operate with one or two CMM modules installed. If there are two CMM modules, one management processor is considered “primary” and is actively managing the system. The other management processor is considered “secondary” and remains ready to quickly take over management in the event of hardware or software failure on the primary. The switch fabric on the CMM operates independently of the management processor. If there are two CMM modules installed, both fabric modules are active. Two CMM modules must be installed in the OS9700 & OS9800 to provide full fabric capacity. If only one CMM module is installed, then there is no management or fabric redundancy and the system capacity is halved.

	In a dual synchronized CMM environment, if a user executes the "takeover" command, then only the processor module of the previous primary CMM would go down and the fabric module would still remain up. This would result in no packet drop. **Note:** The OS9600 will operate with one CMM module only, therefore the "fail-over", the "capacity" and the "takeover" concepts as described above do not apply to the OS9600 chassis.
Image downloadable time to the switch & the approximate size of the AOS Code	Based on the connection speed The size of the AOS Code is approximately: 24-32 MB
System Resiliency Verification	Alcatel.Lucent OmniSwitch 9000 switches are designed in such a way that is highly reliable under extreme stress conditions. The OmniSwitch 9000 switches are rigorously tested to ensure that the system is able to sustain heavy loads and allow for continued availability of all system resources. The typical test setups involve: • Running in normal operational mode where system is running under the specified CPU threshold values on both CMM and NIs. • Running above the CPU threshold values all the time.
Routing Information Base (RIB) & Forwarding Information Base (FIB)	Tested figures: The RIB is 96K (IPv4) while the FIB is 12K (IPv4).
Layer 3 Network Convergence Describe how your equipment implements this along with how the FIB is updated.	Local interface routes and static routes are immediately populated into the FIB on the CMM during the boot process which is then loaded onto each of the Network Interface cards which then installs the FIB into the ASIC hardware forwarding tables. As dynamic routes are learned via the routing protocols these routes are also installed in the FIB and distributed to the Network Interface cards for installation into the ASIC hardware forwarding tables. Whenever the actual FIB is larger than the HW capacity (12K), the AOS software is capable of using the HW capacity as a cache for the most active entries. If this mode, HW based forwarding will only be achieved for the entries present in the HW.
Layer 2 Network Convergence Describe how your equipment implements this along with how the FIB is updated.	The L2 FIB is essentially composed of MAC addresses, ports, and VLANs. The L2 FIB is populated first by configured static MAC entries. Source learning dynamically adds to the L2 FIB adding MAC-port-VLAN records as a function of possible MAC checks, VLAN checks, port checks (depending on what has been configured). Wherever possible these are translated into entries into the corresponding ASIC table. Because of the possibility of loops and there necessary prevention (especially for broadcast and unknown destination traffic), we generally must implement spanning tree which has an effect on convergence - if 802.1w mode is selected (the default on 6.1.2.R03 and later), typical convergence times are < 1 sec.
Layer-2/Layer-3 Switching	
Root bridge priority / path cost:	• Default spanning tree mode is STP (802.1d.) • The bridge priority can be any value between 0 and 65535 for STP and RSTP protocol in the 16-bit mode. By default spanning tree follows the 16-bit path cost. • The bridge priority can only be in multiples of 4096 in the 32-bit mode or in MSTP mode. • MSTP can only operate in 32-bit mode.
Group mobility Rules supported:	• Port • MAC • MAC range • Mobile-Tag • Protocol • IP • IPX • DHCP port • DHCP MAC • DHCP MAC Range • DHCP Generic
Binding rules supported	• Port-Protocol Binding rule • MAC-Port Binding rule • MAC-IP-Port Binding rule
Rule Precedence:	• Mobile Tag • DHCP Mac • DHCP Mac Range • DHCP Port • DHCP Generic • Mac-Port-IP Binding • Mac-Port Binding • Port-Protocol Binding • Mac • Mac Range

	• Network Rule • Protocol
Max. no. of 1x1 STP instances supported per system	253
VLAN	Port based, IEEE 802.1Q VLANs **4.1.11** Advanced VLAN Classification: MAC, protocol, IP subnet IEEE 802.1ad VLAN Stacking (aka Q-in-Q)
Maximum VLANs **4.2.15**	VLAN Range Support Up to 4094 VLANs for Flat Spanning Tree mode/MSTP and 253 VLANs for 1x1 Spanning Tree mode are supported. In addition, it is now possible on the OmniSwitch 6800/6850/9000 to specify a range of VLAN IDs when creating or deleting VLANs and/or configuring VLAN parameters, such as Spanning Tree bridge values.
VLAN Stacking	The IEEE 802.1Q-in-Q VLAN Tagging purpose is to expand the VLAN space by tagging the tagged packets, thus producing a "double-tagged" frame. The expanded VLAN space allows the service provider to provide certain services, such as Internet access on specific VLANs for specific customers, and yet still allows the service provider to provide other types of services for their other customers on other VLANs. Maximum frame size With the insertion of a 4-byte svlan tag by VLAN Stacking, the maximum frame size that can be accommodated is jumbo frame size less 4 bytes = 9216 – 4 = 9212 bytes. Maximum number of SVLANs: •For port level VLAN Stacking: 4093 (VLAN 2 through 4094). •For port / vlan level VLAN Stacking: 768 (can use any number from 2 through 4094 inclusive).
VLAN Tag Translation (aka "VLAN Tag Overlapping")	VLAN Tag Translation is supported.
Maximum number of BPDUs the switch can handle	Approximately 800 BPDUs per second
MAC Address Table	In synchronized mode (default), up to 16K MAC Addresses is supported per system In Distributed mode, up to 64 K MAC Addresses is supported per system (no more than 16K per NI). 1K (authenticated / mobile users) per system Latency: <10µsec
L2 MAC Address Table Size Enhancement AOSv6.1.3r01 Release	There are now two source learning modes available for the OmniSwitch 9000 Series switches: synchronized and distributed. By default the switch runs in the synchronized mode, which allows a total MAC address tables size of 16K per chassis. Enabling the distributed mode for the switch increases the table size to 16K per module and up to 64K per OmniSwitch 9000 chassis. The 6.1.3.R01 release provides support for this feature on the OmniSwitch 9000 Series; increasing the MAC address table size is not supported on the OmniSwitch 6800 Series and OmniSwitch 6850 Series.
IP Address Table Routes	40K routing table 12K forwarding LPM entries, 8K hosts entries per module Latency: <10µsec
Layer-2 Table Hashing	The L2 Table size is 16K entries. This is organized as 2K buckets with each bucket having 8 entries. The search key for the L2 Table is the 60 bit (i.e. 48-bit DA MAC address + 12 bit VLAN-ID) in the Ethernet MAC header in the incoming flow. The key is hashed into a 11-bit value used to select the bucket in the table using a CRC32 lower 11-bits algorithm. Each entry in the selected bucket is compared with the key. The match must be an exact match since if it does, it must be a host MAC address entry. If the key matches an entry in the bucket, then the information in the entry is used in the ingress logic for the destination port
RSTP Performance Sub-second performance	Link Fail-over: 459ms Link Fail-over Reverse: 240ms Port Fail-over: 220ms Port Fail-over Reverse: 140ms AGG Links Fail-over: 958ms AGG Links Fail-over Reverse: 260ms AGG Fail-over: 219ms AGG Fail-over Reverse: 280ms
Max number of configured VLANs per port	1 K (1,024) with support of full 4K IEEE 802.1Q VLAN Spectrum. Port based (w / IEEE 802.1Q) VLANs. The switch has indeed been tested with up to 4,094 active VLANs as well, but this is really based on switch configuration and available resources. Otherwise, the more practical, or more realistic and/or recommended one is the 1,024 active VLANs.
A-VLAN	Maximum number of Avlan authenticated user per system: 1024. The system supports up to 1024 authenticated/mobile Mac-addresses AVLAN supports RADIUS or LDAP as authentication servers. By configuring multiple servers, user can gain server failover in case of server outage. **Supported rules for AVLAN.** MAC-Port Binding rule MAC-IP-Port Binding rule MAC range (used for IP phone OUI Mac-addresses for instance)
Supported rules for AVLAN	• MAC-Port Binding rule

	• MAC-IP-Port Binding rule • MAC range (used for IP phone OUI Mac-addresses for instance)
Max number of configured VLANs per system	4K (4,094) The switch has indeed been tested with up to 4,094 active VLANs, but this is really based on switch configuration and available resources. In the STP flat Mode: 4K VLANs are supported over 802.1Q or over a trunk. In the STP 1x1 Mode: 253 VLANs are supported over 802.1Q or over a trunk. In the STP Multiple Mode (IEEE 802.1s): 4K VLANs amongst 16 Multiple STP Instances (MSTPI).
Max number of system wide Rules	8 K (8,192)
Max number of Link Aggregate **4.2.11**	32 aggregates of up to 8 ports each, across modules **4.2.12** Support for static aggregate (aka OmniChannel) Support for dynamic aggregate (IEEE 802.3ad) ***LOAD BALANCE ALGORITHM*** The load balance is the same for static and LACP link aggregation. The load balance takes the 3 last bits of the source address and the 3 last bits of the destination address and does an XOR. That gives a number between 0 and 7 Note that Link1 is the lowest port number, then Link2 is next port number …
DHCP	DHCP Relay, Option 82 & Snooping (including port-MAC-IP binding)
DHCP Option-82 description	The DHCP relay agent information option (option 82) enables a Dynamic Host Configuration Protocol (DHCP) relay agent to include information about itself when forwarding client-originated DHCP packets to a DHCP server. The DHCP server can use this information to implement IP address or other parameter-assignment policies. The following events will occur when DHCP relay agent option 82 is enabled on the switch. • A DHCP client broadcasts a DHCP request to the network. • Switch (DHCP relay agent) get a copy of the DHCP request, adds relay agent option 82 to the DHCP request packet and then forwards it to the configured DHCP server. • The DHCP server receives the DHCP request packet with the option 82 field. If the DHCP server is option 82 capable, it will assign an IP address based on that option 82 information. Otherwise this option 82 field will be ignored by the DHCP server. • The DHCP server unicasts a DHCP offer with option 82 to the switch. • The switch removes option 82 and forwards it back to the DHCP client. Feature to be supported with AOS 6.1.3R01
Auto-negotiation	Speed (10, 100, 1000Mbps) and duplex mode (half or full)
Traffic Control	IEEE 802.3x (Flow Control) (Note: the switch supports and honors the incoming IEEE 802.3x pause frames, but it does not generate outgoing IEEE 802.3x pause frames)
Spanning Tree	IEEE 802.1D Spanning Tree Protocol (STP) – 1998 / 2004 edition IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) – 2001 edition IEEE 802.1s Multiple Spanning Tree Protocol (MSTP) – 2002 / 2005 edition Support of single and multiple instances for STP & RSTP BPDU Watch Guard How many Multiple Spanning Tree Groups are supported? 253 Is one Spanning Tree per Group supported? Yes only in a 1x1 STP mode Is one Spanning Tree per port supported? Yes Is Single Instance Spanning Tree supported? Yes only in a flat STP mode
Port Monitoring	The same NI cannot support both mirroring and monitoring configuration i.e. a user cannot have a port monitoring and a port mirroring session on the same NI Only one monitoring session at a time across the entire system Only the first 64 bytes of the packet can be monitored. Due to the port monitoring file size, the system can only store the first 2K packets (i.e. 140K/64 = 2187) The port monitoring is not supported on the LinkAgg ports. Enabling the monitoring function affects the performance. Consequently, Port Monitoring performance is not at wire-rate.
Max number of Port Mirroring sessions	The N-to-1 port mirroring allows the user to specify multiple numbers of ports, range of ports as mirrored source in a single command. Maximum number of mirror source ports could be set to 128, this is enhancement in 6.1.3.R01. Aggregate ports are allowed to be mirrored on the physical ports. Mirroring on the logical LinkAgg port ID is not supported. Mirroring Sessions Supported: One session supported per standalone switch
Port Mapping (aka. “Private VLAN”)	Port Mapping is a Layer 2 security feature providing port-based security and isolation between ports within a VLAN. It is an extension of the common VLAN implementation. Port Mapping provides security and isolation between two set of ports (typically referred as “users” and “uplinks” set) on a switch so that traffic from the "users" ports can only be sent to the uplinks and cannot travel to another port within that switch. When Port Mapping is enabled, there is no forwarding of any sort (unicast, broadcast, or multicast) between ports of the “users” set on a switch, and all traffic between ports on the switch must be forwarded through a designated (router) Layer 3 device, connected on the port of the “uplinks” set. Port Mapping enables per port security, requiring only a VLAN on every switch, not

	every port. This feature greatly minimizes the number of VLANs required. Port mapping feature is supported on OS6800/6850/9000. Following are the limitations for the feature. • 8 sessions supported per standalone switch and stack • An aggregable port of a link aggregation group cannot be a mapped port and vice versa • A mirrored port cannot be a mapped port and vice versa • A mobile port cannot be configured as a network port of a mapping session
STP convergence time (flat, 1x1, 802.1s)	30 sec
802.1w rapid reconfiguration	Less than 1 sec
Learned MAC addresses per port	Up to 16 K MAC Addresses is supported
Learned MAC addresses per system	Up to 64 K MAC Addresses is supported In synchronized mode (default), up to 16K MAC Addresses is supported per system In Distributed mode, up to 64 K MAC Addresses is supported per system (no more than 16K per NI).
Layer-2 forwarding on Ethernet ports	Wire-speed (64 Bytes packets)
Layer-2 forwarding GigE, known MAC	Wire-speed (64 Bytes packets)
Broadcast per Ingress port	Programmable
Loopback Interface	The loop-back interface allows you to uniquely identify a router in the network with one IP address. The advantage of the loop-back interface is to be independent of the physical ip interfaces. In a redundant routing network, the loop-back interface is always accessible when routing topology changes or ip interfaces go down. The main advantage of Loop-back interface is a more reliable Network Management path through OmniVista or an NMS station. Also, you can use the loop-back interface to uniquely identify the router within OSPF and BGP if you set the router-id to the same as the loop-back address. The loop-back can also be used for the RP (Rendezvous Point) in PIM-SM. The loop-back address is also used for the sFlow Agent IP address. The Loopback address is used for source IP of RADIUS authentication.
User Definable Loopback Interface	Loopback0 is the name assigned to an IP interface to identify a consistent address for network management purposes (including SNMP/sFlow datagrams). The Loopback0 interface is not bound to any VLAN; therefore it always remains operationally active. This differs from other IP interfaces, such that if there are no active ports in the VLAN, all IP interfaces associated with that VLAN are not active. In addition, the Loopback0 interface provides a unique IP address for the switch that is easily identifiable to network management applications.
Sever Load Balancing (SLB)	There are 2 kind of server clusters: -Server Farm: The traffic is truly destined to the Server Farm and the destination IP is the Virtual IP of the Server Farm. Each server is also configured with a Loopback Interface for the Virtual IP -Advanced Clustering: the traffic is not necessarily destined to a Virtual IP, instead, it is matching a user defined QoS condition, allowing L1-L4 classification. The most common application is Firewall clustering where packets are load balanced to several firewall for inspection and sent back if accepted. **The following values are the tested limits with the functionality verified (stress test):** Tested limit of clusters (on a per switch basis) is 16. Tested limit of servers (on a per cluster basis) is 16. Tested limit of Probes: 20 Probes **The following values are the maximum limits enforced by the Code:** Maximum number of clusters: 16 Maximum number of physical servers: 75 Maximum number of probes on a switch: 20 Sever Load Balancing (SLB) Health monitoring is performed by the CPU of the Primary Management. LOAD BALANCING HASHING In both “VIP” and “Condition” SLB, the traffic is balanced among the servers using an hash algorithm based on IPSA and IPDA. Internally, each active server is seen as a host ECMP route to reach the cluster. Therefore, the load balancing is the same than the ECMP load balancing.
Layer-3 Routing Unicast (IPv4)	
Large L3 table support	**Hardware:** • Maximum number of active flows in the hardware: 12K One active flow is usually one “remote-subnet” flow (not a per destination ip flow based) Now with the ARP table enhancement, one active flow can also be a “host routed” flow The table is shared for - IPV4 active flow (remote ipv4 network): 1 entry - IPV6 active flow (remote ipv6 network): 2 entries - Host active flow (ARP entry): 1 entry • Maximum number of active “ARP entries” flows: 8K • Maximum number of ECMP Next-hops that can be stored: 512 **Software:**

	• Maximum number of IPv4 routes that can be held in the software routing table: 96K • Maximum number of IPv6 routes that can be held in the software routing table: 5K • Maximum number of ARP entries that can be held in software ARP table: 16K
IP Routing	Static Routing, RIPv1&v2, OSPFv2, and BGPv4 (including graceful restart) **4.2.13**
Maximum number of IP route entries (Layer-3 Routing Table Size) (Maximum Routing Information Base – RIB)	Up to 96K routing table is supported. 12K forwarding LPM entries, 8K hosts entries per module. Latency: <10µsec
Max number of IP Router interfaces per system – Single mode	1 K (1,024)
Max number of IP routes	Up to 96K
Max number of IP static routes	1 K (1,024) routes
RIPv1&v2	**The following values are the maximum limits enforced by the code.** • Maximum number of RIPv2 interfaces per router: 2K • Maximum number of RIPv2 routes: Unlimited **The following values are the tested limits with the functionally verified (stress test).** • Tested number of RIPv2 interfaces per router: 10 • Tested number of RIPv2 peers per router, one per interface: 10 • Tested number of RIPv2 routes with no redistribution from OSPFv2 RIB: 8500
OSPFv2 Specifications	**The following values are the maximum limits enforced by the code.** Maximum number of Areas (per router): 32 Maximum number of Interfaces (per area): 100 Maximum number of Interfaces (per router): 32 x 100 (Limited only by max. num of IPv4 interfaces = 4096) Maximum number of Link State Database entries (per router): Unlimited Maximum number of neighbors/adjacencies (per router): 254 Maximum number of neighbors/adjacencies (per area): 128 Maximum number of routes (per router): Unlimited Maximum number of OSPFv2- ECMP gateways (per destination): 4 Max number of OSPFv2 Sessions: 1 **The following values are the tested limits with the functionally verified (stress test).** **On OS9000 ABR routers:** Tested number of IP Routers on OS9000 router: 32K Tested number of OSPFv2 Routes on OS9000 router: 32K Tested number of OSPFv2 Interfaces on OS9000 ABR: 128 Tested number of OSPFv2 Areas on OS9000 ABR: 6 Tested number of OSPFv2 Adjacencies on OS9000 ABR: 128 Tested number of LSAs on OS9000 ABR: 32K Tested number of OSPFv2- ECMP gateways (per destination): 4 Tested number of OSPFv2 Sessions: 1 **On OS9000/OS6850 non-ABR routers:** Tested number of IP Routers on OS9000/OS6850 router: 96K Tested number of OSPFv2 Routes on OS9000/OS6850 router: 96K Tested number of OSPFv2 Interfaces on OS9000/OS6850 ABR: 27 Tested number of OSPFv2 Areas on OS9000/OS6850 ABR: 6 Tested number of OSPFv2 Adjacencies on OS9000/OS6850 ABR: 27 Tested number of LSAs on OS9000/OS6850 ABR: 24K Tested number of OSPFv2- ECMP gateways (per destination): 4 Tested number of OSPFv2 Sessions: 1
ECMP	Only 512 networks can be programmed in the ECMP table, so that the flows can be load balanced among the different paths. When having more than 512 ECMP routes on the “show ip route”, only the last (highest) 512 routes are programmed in the ECMP table. • Only 512 networks can be load balanced over ECMP links • The other “ECMP networks” will always be routed on the same link (single path used).
BGP Routing Limitations	**The following values are the maximum limits enforced by the code.** Maximum BGP Peers per Router: 32 Maximum number of routes supported: Unlimited Range for AS Numbers 1 to 65535 Range of Local Preference Values 0 to 4294967295 Range for Confederation IDs 0 to 65535 Range for MED Attribute 0 to 4294967295 **The following values are the tested limits with the functionally verified (stress test).** Tested BGP Peers per Router: 32 Tested number of routes supported: 65,000 Range for AS Numbers 1 to 65535 Range of Local Preference Values 0 to 4294967295 Range for Confederation IDs 0 to 65535

	Range for MED Attribute 0 to 4294967295
ARP Table: Max number of ARP entries per system	Up to 8K L3 ARP entries are supported.
Layer-3 forwarding, known IP@64 bytes pkt	Wire-speed
Layer-3 forwarding, known IP@1518 bytes pkt	Wire-speed
Layer-3 forwarding, known IP@ Jumbo pkt	Wire-speed
Trunking 2 VLANs, 64 Bytes pkt	Wire-speed
Trunking 2 VLANs, 1518 Bytes pkt	Wire-speed
RIP Learning Rate	500 / sec
OSPF Learning Rate	500 / sec
Route Convergence for OSPF	1.2 sec
IPv4 redistribution	Supported platform: OS6800, OS6850, and OS9000 IPv4 Redistribution instances use route-maps to redistribute routes from a source protocol RIB to a destination protocol RIB. The source protocol can be BGP, RIP, OSPF, Local or Static. The destination protocol can be BGP, RIP or OSPF. **The following values are the tested limits with the functionally verified (stress test).** • Tested number of route-maps that can be created on router: 200 • Tested number of route-map sequences that can be created on router: 400 • Tested number of IPv4 access-lists that can be configured on router: 200 • Tested number of OSPFv2 routes that can be redistributed into RIPv2: 8.5K • Tested number of RIPv2 routes that can be redistributed into OSPFv2: 8K
Multicast & Network Protocols & Resilience	
Groups	1 K groups
Multicast support **4.2.14**	IGMPv1&v2&v3 Snooping MLD Snooping (IPv6) DVMRP PIM-SM PIM-DM
Flow Table	1021 entries per system
VLAN Replication	2048 entries per system
Max number of DVMRP Interfaces	128
Max number of DVMRP Neighbors	256
Max number of DVMRP Tunnels	1 per interface
Max number of PIM-SM Interfaces	128
PIM-DM (IPv4) Note: IPv6 PIM-DM will be supported in a future Release	PIM-DM is a multicast routing protocol that defines a multicast routing algorithm for multicast groups that are densely distributed across a network. It uses the underlying unicast routing information base to flood multicast datagrams to all multicast routers. Prune messages are used to prevent future messages from propagating to routers with no group membership information. It employs the same packet formats as sparse mode PIM (PIM-SM). PIM-DM assumes that when a multicast source starts sending, all downstream systems want to receive multicast datagrams. Initially, multicast datagrams are flooded to all areas of the network. PIM-DM uses RPF (Reverse Path Forwarding) to prevent looping of multicast datagrams while flooding. If some areas of the network do not have group members, PIM-DM will prune off the forwarding branch by instantiating prune state. PIM-DM differs from PIM-SM in two essential ways: 1. There are no periodic joins transmitted, only explicitly triggered prunes and grafts. 2. There is no Rendezvous Point (RP). This is particularly important in networks that cannot tolerate a single point of failure.
IGMP learning performance	The system can process 1000 IGMP per second. However, the performance can drop to 128 when IGMP are received too fast. • Burst of 1000 IGMP reports at 1000 packet/sec: all 1000 groups are learnt • Burst of 1000 IGMP reports at 1Gbps: only 128 groups are learnt
Zapping	You can configure “ip multicast zapping” to optimize channel surfing. That will instantly stop forwarding multicast to a client when that client sent an IGMP Leave. The zapping time can be measured by the leave message received by the switch and the last packet received by the client. This is usually in milliseconds. The feature is well suited for Multicast Switching and zapping only works well when “ip multicast querying” is disabled.
L2 static multicast	•1022 static multicast MACs are supported on OS6850 and OS9000. The L2 Multicast table can have 1024 entries but 2 are reserved for other applications.
Multicast without 8021.Q on 10/100Mbps interfaces	Wire-speed
Multicast without 8021.Q on 1000Mbps interfaces	Wire-speed
Multicast with 8021.Q, 0 copies, 1518Bytes pkt on 10/100/1000Mbps ports and/or GigE ports	Wire-speed
Multicast with 8021.Q, 1 copies, 1518Bytes pkt on 10/100/1000Mbps ports and/or GigE ports	Wire-speed
Multicast with 8021.Q, 2 copies, 1518Bytes pkt on 10/100/1000Mbps ports and/or GigE ports	Wire-speed

Network Protocols	Generic UDP Relay (including DHCP Relay) TCP/IP Stack NDP ARP
Resilience	
VRRPv3	Virtual Router Redundancy Protocol, VRRPv3, is designed to eliminate the single point of failure existing in a static default routed IPv6 environment. The loss of the default router isolates all systems not able to detect an alternate path. VRRPv3 provides the capability for assigning the responsibility of a virtual router to one of the IPv6 VRRPv3 routers on a LAN. A total of 255 VRRP3 instances can be configured if only IPv6 instances are configured. The total of 255 instances on a box is the maximum number of VRRP instances (VRRP2 + VRRP3) that can be configured on a box.. As an example if a user configures 200 VRRP2 instances, then only 55 VRRP3 instances can be configured. If a user configures 255 VRRP2 instances then no VRRP3 instances can be configured and vice versa.
Layer-3 Routing Unicast (IPv6)	
Large L3 table support	Note: ARP is referring to a 32bit entry associated with IPv4. For IPv6, 128bits, we are talking of NDP (equivalent to ARP for IPv6) **Hardware:** • Maximum number of active flows in the hardware: 6K One active flow is usually one "remote-subnet" flow (not a per destination ip flow based) Now with the NDP table enhancement, one active flow can also be a "host routed" flow The table is shared for - IPV4 active flow (remote ipv4 network): 1 entry - IPV6 active flow (remote ipv6 network): 2 entries - Host active flow (NDP entry): 2 entries • Maximum number of active "NDP entries" flows: 8K • Maximum number of ECMP Next-hops that can be stored: 512 **Software:** • Maximum number of IPv4 routes that can be held in the software routing table: 96K • Maximum number of IPv6 routes that can be held in the software routing table: 5K • Maximum number of NDP entries that can be held in software NDP table: 16K
IP Routing	Static Routing, RIPng, OSPFv3, and Multiprotocol Extensions for BGPv4
Maximum number of IP route entries (Layer-3 Routing Table Size) (Maximum Routing Information Base – RIB)	Up to 16K routing table is supported. 6K forwarding LPM entries, 4K hosts entries per module. Latency: <10µsec
Max number of IP Router interfaces per system – Single mode	1 K (1,024)
IPv6 routes	The total number of IPv6 routes supported in hardware (with no IPv4 routes) is 6000.
Max number of IPv6 static routes	1 K (1,000) routes
IPv6 routing interfaces	The recommended number of IPv6 routing interfaces is 100
IPv6 prefixes per routing interface	The recommended number of IPv6 prefixes per routing interface is 50
IPv6 global unicast addresses per routing interface	The recommended number of IPv6 global unicast addresses per routing interface is 50
A 6to4 tunnel	A 6to4 tunnel explicitly uses an "ingress tunnel" for each IPv4 interface configured on the system. The limit is 100 ingress tunnels The 10GIG routing performance over an IPv6 tunnel (6to4 and configured tunnel) has been determined to be 10,775,862 – 96 byte packets per second. The 10GIG routing performance NI-NI or Single NI has been determined to be 14,880,812 - 64 byte packets per second.
RIPng	**The following values are the maximum limits enforced by the code.** The total number of RIPng interfaces is 100. The maximum number of RIPng neighbors is 20 Maximum number of RIPng routes: 5K routes (Depending on the number of RIPng interfaces, and neighbors configured the maximum number of routes may vary.) **The following values are the tested limits with the functionally verified (stress test).** (a) Tested number of RIPng interfaces per router: 10 (b) Tested number of RIPng peers per OS9000 router: 10 (c) Tested number of RIPng routes with no redistribution from OSPFv3 RIB: 1000
NDP Table: Max number of NDP entries per system	Up to 8K (8,192) L3 NDP (ARP) entries are supported. Note that ARP (IPv4) and NDP (IPv6) are using the same resources.
OSPFv3 Specifications	**The following values are the maximum limits enforced by the code.** Maximum number of Areas (per router): 32 Maximum number of Interfaces (per router): Unlimited (Limited only by max. num of IPv4 interfaces = 4096) Maximum number of Interfaces (per area): 100 Maximum number of Link State Database entries (per router): Unlimited

	Maximum number of adjacencies (per router): adjacency is no different from neighbor, below. Maximum number of OSPF- ECMP gateways (per destination): 4 Maximum number of neighbors (per router); 254 Maximum number of neighbors (per area); 64 Maximum number of routes (per router): Unlimited **(Future Release)** (Depending on the number of Areas, Interfaces, Adjacencies, and Neighbors configured, the maximum number of routes may vary.) Max number of OSPF Sessions: 1 **The following values are the tested limits with the functionally verified (stress test).** **On an OS9000 ABR Routers:** (a) Tested usable Hello Interval with 20 Interfaces in 5 Areas with 4 Interfaces in each Area: 5 sec (b) Tested usable Router Dead Interval with 20 Neighbors, 4 each in 1 Area for 5 Areas: 20 sec (c) Tested usable number of LSAs that the OS9000 router can stably hold: 5K (d) Tested usable no. of Ospfv3 Routes that the OS9000 router can stably hold in this scenario: 5K (e) Tested m number of usable Ospfv3 Interfaces in 5 areas with 5K LSAs: 20 (f) Tested number of usable Ospfv3 Neighbors in 5 areas with 5K LSAs: 20 (g) Tested number of usable Ospfv3 Interfaces between any two OS9000 routers: 4 (h) Tested number of usable Ospfv3 Areas on a OS9000 ABR: 5 (i) Tested number of OSPF Interfaces on OS9000 ABR: 20 (k) Tested number of OSPF Areas on OS9000 ABR: 5 (l) Tested number of OSPF Adjacencies on OS9000 ABR: 20 (m) Tested number of LSAs on OS900 ABR: 5K **On OS9000 non-ABR routers:** Numbers for an OS9000 non-ABR will be a sub-set of the above numbers for an OS9000 ABR. Tested usable Hello Interval with 20 Interfaces in 5 Areas with 4 Interfaces in each Area: 5 sec Tested usable Router Dead Interval with 20 Neighbors, 4 each in 1 Area for a total of 5 Areas: 20 sec Tested number of IP Routes on OS9000 router: 5K Tested number of OSPFv3 Routes on OS9000 router: 5K Tested number of usable Ospfv3 Interfaces between any two OS9000 routers: 4 Tested usable number of LSAs that the OS9000 router can stably hold: 5K
Layer-3 forwarding, known IP@64 bytes pkt	Wire-speed
Layer-3 forwarding, known IP@1518 bytes pkt	Wire-speed
Layer-3 forwarding, known IP@ Jumbo pkt	Wire-speed
Trunking 2 VLANs, 64 Bytes pkt	Wire-speed
Trunking 2 VLANs, 1518 Bytes pkt	Wire-speed
RIP Learning Rate	500 / sec
OSPF Learning Rate	500 / sec
Route Convergence for OSPF	1.2 sec
IPv6 REDISTRIBUTION	**The following values are the tested limits with the functionally verified (stress test).** (a) Tested number of route-maps that can be created on an OS9000 router: 200 (b) Tested number of route-map sequences that can be created on an OS9000 router: 400 (c) Tested number of IPv6 access-lists that can be configured on an OS9000 router: 100 (d) Tested number of OSPFv3 routes that can be redistributed into RIPng: 1K (e) Tested number of RIPng routes that can be redistributed into OSPFv3: 1K
Multinetting	
Multinetting This feature allows IP traffic from multiple subnets to coexist on the same VLAN. A network is said to be multinetted when multiple IP subnets are brought together within a single broadcast domain (VLAN). It is possible to assign up to eight different IP interfaces per VLAN.	A network is said to be multinetted when multiple IP subnets are brought together within a single VLAN. For example, one may configure the subnet 192.168.1.0/24 and 194.2.10.0/24 to run on the same switch interface. In other words, traffic from the 192.168.1.0 subnet and traffic from the 194.2.10.0 subnet would coexist on the same physical VLAN. Within a Layer 2 environment, the traffic is broadcast between all subnets configured in the same VLAN. Layer-3 traffic is routed between the configured subnets in the same VLAN. Possible uses for Multinetting: • Subnet renumbering – used during transition from one addressing scheme to another to maintain connectivity. • Ability to support more hosts on one physical link – used to add more hosts to a broadcast domain than the addressing scheme allows. • Supporting multiple subnets on one interface where configurations do not allow complete separation of subnet traffic. For example, a college campus may have departments where users are connected to a switch via hubs. Connected to each of the hubs are users configured to be in different subnets. The hubs are connected to the switches using port-based vlan configuration. Network administrators use Multinetting so they do not have to worry about re-cabling or reconfiguring ports for users in different subnets.
Supported features:	• Up to 8 subnets per VLAN • All existing dynamic routing protocols, routing between each of the multinetted subnets in one VLAN and routing between each of the multinetted subnets and other VLANs • VRRP DHCP is only supported on the primary interface of the multinetted vlan. All devices are assigned to

	the same scope (the one for the primary interface) With VRRP and Multinetting, you can still configure multiple instances to load balance the master role among the sub-netted interfaces.
Routing In Multinetting	Routing protocols (RIP, OSPF, and BGP) are supported in a multinetted environment. The routing interfaces are now based on ip interfaces, instead of the VLANs. Therefore, routing protocols are totally independent of VLANs and their data structures are maintained as part of an array indexed by ip interface only. There is no difference between running a routing protocol on an interface part of a multinetted vlan or a regular interface. Each subnet (interface) on the multinetted vlan can run its own routing protocol.
Multicast Routing In Multinetting	The multicast routing protocols will be supported on one interface per VLAN. One interface designated the primary interface, will be used for the multicast routing protocols. The multicast routing protocols will not allow configuration on any non-primary interfaces. By default the first interface is the primary interface. DVMRP and PIM-SM will only allow configuration on the primary interface of a VLAN. This is to ensure consistency between the multicast routing protocols (DVMRP, PIM-SM, IPMRM), IPMS and IGMP.
Layer-3 Routing (IPX)	
Routes	1K Routes 1K Host entries
IPX Routing	64 IPX interfaces Static routing (256 routes) RIP/SAP, 1K routes 5000 RIP and SAP entries each are supported. IPX routing is limited to 8000 packets per second per NI. Each NI can independently route up to 8000 p/s.
Policy/QoS	
QoS / ACLs	Features summary: • 802.1p classification **4.1.8** • TOS/DSCP classification **4.1.9** • Ethertype classification • IP protocol classification • ICMP type and code classification • TCP Flag classification and “established” for implicit “reflexive” tcp flows • “qos apply” will not impact existing flows • Port disable rules to shutdown a port when incoming packets matches a rule • Rule logging • Port redirect action to force a packet to be sent out on a given port • User port profiles to filter and shutdown ports for BPDUs, IP spoofing and routing protocols (RIP, OSPF, BGP) • DropServices to drop tcp/udp ports • IGMP ACLs • L2/L3/L4 QoS fully supports IP multicast traffic (priority, bandwidth shaping..) • 8 hardware queues per port
QoS Conditions & Actions supported	The following types of conditions are available: • L1 conditions: source port, destination port, source port group, destination port group • L2 conditions: source Mac, source Mac group, destination Mac, destination Mac group, 802.1p, ethertype, and source vlan (Destination vlan is not supported). • L3 conditions: ip protocol, source ip, source network group, destination ip, destination network group, TOS, DSCP, ICMP type, ICMP code. • L4 conditions: source TCP/UDP port, source TCP/UDP port range, destination TCP/UDP port, destination TCP/UDP port range, service, service group, tcp flags • IP multicast condition: An ip multicast condition is used for IGMP ACLs. The multicast ip is actually the multicast group address used in the IGMP report packet. IP multicast can be combined with destination port, destination vlan, destination Mac, destination ip, that are the port/vlan/Mac/ip of the device that sent the IGMP report The following actions are available: • ACL (disposition drop/accept – default is accept) • Priority • 802.1p/TOS/DSCP Stamping • 802.1p/TOS/DSCP Mapping • Maximum bandwidth • Redirect Port Note: Condition combinations and Action combinations are also supported.
Priority Queues	Eight hardware based queues per port
Traffic Prioritization	Flow based QoS in hardware (L1-L4) Internal & External (aka remarking) prioritization

<table>
<tr><td>Traffic redirection</td><td>Policy-based routing
Server load balancing (including health monitoring of servers)</td></tr>
<tr><td>Bandwidth Management</td><td>Flow Based bandwidth management, ingress policing / egress shaping, 64kbps granularity
Port based egress shaping, 1Mbps granularity</td></tr>
<tr><td>Queue Management</td><td>Configurable de-queuing algorithm
Strict Priority
Weighted Round Robin
DRR (Deficit Round Robin). This mode is quite similar as WRR
In the Strict Priority mode, a port has 8 strict priority queues (SPQ) and all the queues on the port are serviced strictly by priority.
In the WRR or DRR, queues are serviced on a round robin based on their weight. The higher the queue weight, the higher is the throughput for that queue. Any queue can be configured with a weight of 0 to make that queue strict priority. The weight ordering does not need to follow the queue order.</td></tr>
<tr><td>Queuing Scheme and Servicing Mode</td><td>OS9000 has 8 queues per egress port
OS9000 has 4 Queuing schemes per egress port:
Strict-Priority (default mode)
WRR (Weighted Round Robin)
DRR (Deficit Round Robin). This mode is quite similar as WRR
Priority-WRR. Mixed of strict priority and WRR
In the Strict Priority mode, a port has 8 strict priority queues (SPQ) and all the queues on the port are serviced strictly by priority.
In the WRR or DRR, queues are serviced on a round robin based on their weight. The higher the queue weight, the higher is the throughput for that queue. Any queue can be configured with a weight of 0 to make that queue strict priority. The weight ordering does not need to follow the queue order.</td></tr>
<tr><td>Queue Mapping Table</td><td>Queue Mapping Table
<table>
<tr><td>802.1p</td><td>TOS / DSCP</td><td>Priority Rule</td><td>Egress Queue</td><td>Servicing</td></tr>
<tr><td>0</td><td>0 / 0-7</td><td>0</td><td>0</td><td>SPQ or W</td></tr>
<tr><td>1</td><td>1 / 8-15</td><td>1</td><td>1</td><td>SPQ or W</td></tr>
<tr><td>2</td><td>2 / 16-23</td><td>2</td><td>2</td><td>SPQ or W</td></tr>
<tr><td>3</td><td>3 / 24-31</td><td>3</td><td>3</td><td>SPQ or W</td></tr>
<tr><td>4</td><td>4 / 32-39</td><td>4</td><td>4</td><td>SPQ or W</td></tr>
<tr><td>5</td><td>5 / 40-47</td><td>5</td><td>5</td><td>SPQ or W</td></tr>
<tr><td>6</td><td>6 / 48-55</td><td>6</td><td>6</td><td>SPQ or W</td></tr>
<tr><td>7</td><td>7 / 56-63</td><td>7</td><td>7</td><td>SPQ or W</td></tr>
</table>
(*) SPQ Strict Priority Queue or Weighted Fair Queue if configured with a weight > 0</td></tr>
<tr><td>Congestion Avoidance</td><td>The congestion avoidance mechanism that is currently supported is built-in the hardware ASIC and managed through our configurable queuing & de-queuing schemes.</td></tr>
<tr><td>Power over Ethernet</td><td>IEEE 802.3af (requires OS9-GNI-P24 & IP-Shelf)
Maximum power of 2400watts (600watts per PSU)</td></tr>
<tr><td>Max number of Rules</td><td>128 per port; 2048 policy rules per chassis</td></tr>
<tr><td>Max number of Actions</td><td>128 per port; 2048 policy actions per chassis</td></tr>
<tr><td>Max number of Conditions</td><td>128 per port; 2048 policy Conditions per chassis</td></tr>
<tr><td>Max number of Policy Services</td><td>256</td></tr>
<tr><td>Max number of Policy Groups</td><td>1024
512 entries per policy group</td></tr>
<tr><td>Max number of Queues</td><td>8 / port</td></tr>
<tr><td>Filtering or ACL Throughput</td><td>Wire-speed</td></tr>
<tr><td>Rule logging</td><td>OS9000 can log the packets matching a policy rule.
The most common use of that feature is to log packet matching an ACL drop policy. To enable logging configure the policy rule with “log [log interval x]”
The log interval is optional and the default interval is 30 sec.
You can configure a log interval between 1 and 3600 sec.
Depending on the configured log interval, the system periodically set the hardware to send copy of the packet matching the rule to CPU. As soon as the CPU receives a packet matching the rule, the system reset the hardware to no longer send copy to CPU until the next interval, to keep CPU low.
The first packet is always logged. If one packet matching the rule is seen during the log interval time, it will be logged.
Limitation:
More than one packet can be logged depending on the rate of the traffic (because of time required by the CPU to stop the sampling).
Log interval less than 5 sec will be accepted by CLI, but logging will be done every 5 sec
Logging does not lot all matching packets (not an IDS)
Note: CPU stays low with rule logging enable. We tested a logging drop rule with 10Gbps of incoming traffic and CPU stays low.</td></tr>
<tr><td>Egress Bandwidth Shaping</td><td><u>Port shaping</u></td></tr>
</table>

	Shaping limits the bandwidth on the egress port. Shaping implies that the shaping function controls the rate at which the egress port sends the packets, regardless of egress queues. The granularity is 64Kbps. **Queue shaping** You can also configure maximum and minimum bandwidth on a per egress queue basis. Configuring an egress queue max bandwidth will shape priority traffic mapped to that queue. Configuring an egress queue min bandwidth will guarantee that bandwidth for priority traffic mapped to that queue. When a queue has a minimum bandwidth configured, traffic within that bandwidth has the HIGHEST priority, regardless the servicing mode or the priority of that queue. **Limitation:** The egress bandwidth shaping is only on a per port basis; the system cannot do a per flow basis egress bandwidth shaping.
Ingress Max Bandwidth Policing	Using policy rule with maximum bandwidth action, you can limit the bandwidth on the ingress. Policing implies dropping the traffic when the programmed rate is exceeded. Policing is on a per flow basis. The granularity is 64kpbs. You can do the following: • Ingress port rate limiting by configure a policy using a source port • Ingress flow based rate limiting by configure a policy defining that flow • Mixed of ingress and flow based rate limiting **Limitations:** • Ingress rate limiting is done at the ingress NI. Policies spread out on multiple NIs will make the total egressing rate to be higher than the configured value (up to the N time the limit where N is the number of NI being spread) • "Show active policy rule" will count the packets that exceed the rate limiting, not the packets that matches the rule
VLAN Ingress Filtering to prevent VLAN leakage	VLAN Ingress Filtering to prevent VLAN leakage is supported only the VLAN(s) statically assigned to a port will be accepted.
VLAN CoS preservation	If the ingress and egress port are both tagged and the ingress port is set to trusted then priority of the frame is enforced and preserved.
VLAN CoS differentiation	Each port provides 8 HW based priority queues, in order to provide VLAN CoS differentiation, it is required to reserve one of these queues to the corresponding VLAN (max of 8 VLAN per port with CoS differentiation).
Untrusted Ports and Packet Priority	On untrusted ports the priority/queue of the incoming packet is based on the "port default 802.1p value". By default, the port default 802.1p value is 0 making traffic to be mapped to Q0 (best effort). Also, regardless or bridging or routing: • 802.1p within the packets is set to the port default 802.1p • DSCP within the packets is set to the port default dscp Changing the port default 802.1p will: • Change the priority of all traffic from that port. That is like a "port priority" • Set the 802.1p value in the packet to that port default 802.1p Changing the port default DSCP will: • NOT change the internal priority • Set the DSCP value in the packet to that of the port default DSCP Notes: On untrusted port, the default 802.1p defines the default internal priority for all packets. Untagged packets on untrusted ports get an 802.1p value from the port default 802.1p (if going out on tagged interface). **Limitation:** On untrusted ports, if the packet matches a policy rule, the DSCP in the packet is unchanged; it is not set to the port default dscp
Trusted Ports and Packet Priority	On trusted ports the priority/queue of the incoming packet is based on the ingress packet 802.1p or ToS/DSCP value. • Non IP packets are prioritized based on the packet 802.1p value • IP packets are prioritized based on the packet TOS/DSCP value Port default 802.1p or DSCP has no effect on trusted ports. Notes: On IP packets, the 802.1p is set to match the packet ToS value. Untagged non-IP packets always get an 802.1p of 0 and priority 0 (if going out on tagged interface). The port default 802.1p is not applied.
802.1p/TOS/DSCP Stamping/Mapping policies	Regardless the condition or classification, the following stamping/mapping actions are allowed • Stamp 802.1p • Stamp TOS (precedence) • Stamp DSCP • Stamp 802.1p and TOS/DSCP • Map 802.1p to 802.1p • Map 802.1p to TOS

	• Map 802.1p to DSCP • Map ToS to 802.1p • Map ToS to TOS • Map ToS to DSCP • Map DSCP to 802.1p • Map DSCP to TOS • Map DSCP to DSCP Stamping/mapping policies change the internal priority of the packets: • Internal Priority is always based on the new 802.1p or TOS/DSCP being stamped/mapped • Stamp/map TOS/DSCP also gives internal priority for non IP packets matching the rule • Mapping rules takes one TCAM rule entry for each entry in the map group • If both 802.1p and TOS/DSCP are stamped in a policy rule, priority is based on the stamped 802.1p value Notes: On trusted ports, stamping/mapping a tos/dscp also change the 802.1p value in the packet to the packet ToS value. If the policy rule has both a 802.1p stamp/map action and a priority action, the packet priority comes from the stamped/mapped 802.1p value, not the priority action.
Policy Rules with Multiple Actions	Multiple policy actions can be combined together within a single rule. The policy actions that can be combined in the same rule are: • Priority • Stamping/mapping • Max BW • Redirect Port
QoS Precedence with Multiple Policy Rules	A flow can match multiple rules but ONLY the action for the highest precedence-matching rule is then enforced. When rule are configured without precedence (default precedence is 0), the first created rule has the highest precedence.
IPv6 Classification & Combinations	Classification & Combinations The following classification criteria are available (in Release 6.1.3.r01) for ipv6 packets • source ipv6 address • destination ipv6 address • Next header. Policies specifying the NH parameter, classify based on the first NH value present in the V6 header of the IPV6 packet • Flow label • TCP Flags/Established. Policies specifying “established” or “tcpflags”, expect the first NH value present in the V6 header to be 6 • ToS/DSCP • source vlan • 802.1p • source Mac • destination Mac • source port • destination port (only for bridged traffic) • Multicast ipv6 for MLD report filtering (similar to IGMP filtering)
IPv6 Actions	Actions All actions are available for Ipv6 policies • ACL (disposition drop/accept – default is accept) • Priority • 802.1p/TOS/DSCP Stamping • 802.1p/TOS/DSCP Mapping • Maximum bandwidth/depth • Redirect Port / Link aggregation
Security	
Switch accessibility under DoS Attack	The following type of packets are processed in software and will increase the CPU usage: • Unresolved L3 packet: unknown destination IP on a local subnet • Broadcast L2 packet (including ARP requests): • IP multicast packet on range 224.0.0.0-224.0.0.255: that includes routing protocol packets such as OSPF, RIPv2 and VRRP packets • All IP packets going to a switch ip interfaces: ping, telnet, http Under normal conditions, the protocol packets are always prioritized in order to maintain the network topology. The following protocol packets are by default prioritized: • BPDUs • OSPF, RIPv2 • VRRP

	• IP multicast protocol (IGMP...) • ARP (both request and reply) **ARP** To prevent an ARP attack, the system limits at 500 pps the number of arp packets sent to CPU (flooding of arp on the network is not limited). Also, there is an early arp discard mechanism to prevent the CPU from processing arp request not destined to a switch ip address. However, under attacks towards the switch, the CPU usage could rise dramatically and makes the switch unreachable for management (WebView, OmniVista or Telnet). In order to keep the switch reachable under attacks, some policies can be created to protect the management access.
Denial of Services (DOS) attacks	The system sustained Denial of Services attacks from Nessus and no switch anomalies (crash or service interruptions) were observed while running the attacks. Nessus has reported the following vulnerabilities: •alya.cgi (Backdoors) •AnalogX denial of service (Denial of Service) •cisco http DoS (Denial of Service) •AnalogX denial of service by long CGI name (Denial of Service) •Jigsaw webserver MS/DOS device DoS (Denial of Service) •Trend Micro OfficeScan Denial of service (Denial of Service) •BadBlue invalid GET DoS (Denial of Service) •DCShop exposes sensitive files (General) •OpenSSH < 3.0.1 (Gain a shell remotely) •Quicktime/Darwin Remote Admin Exploit (Gain a shell remotely) •OpenSSL overflow via invalid certificate passing (Gain a shell remotely) •TESO in.telnetd buffer overflow (Gain root remotely) •OpenSSH AFS/Kerberos ticket/token passing (Gain root remotely) •OpenSSH <= 3.3 (Gain root remotely) •OpenSSH < 3.7.1 (Gain root remotely) •Oracle Application Server Overflow (Gain Root Remotely) •AliBaba path climbing (Remote file access) The following are the versions of Nessus and the Linux platform used. Nessus version: 2.2.0 Linux OS: Fedora Core Release 1
IP security enhancement	Supported platform: OS6800, OS6850, and OS9000 ▪ Detect ARP Flood ▪ Detect packets received with invalid Source IP addresses ▪ Detect packets received with invalid Destination IP addresses ▪ Detect multicast packets with a source MAC that is multicast ▪ Detect multicast packets with mismatching destination IP and MAC address ▪ Detect multicast packets with a Unicast destination IP and Multicast destination MAC address ▪ Detect ping overload ▪ Detect packets with Loopback source IP address
Traffic Filtering	Flow based filtering in hardware (L1-L4)
User Authentication	IEEE 802.1x, with Group Mobility & Guest VLAN* support MAC based Authentication for non-802.1x host Authenticated VLAN (web & telnet based authentication)
Switch protocol security	MD5 for RIPv2, OSPFv2 and SNMPv3 SSHv2 for secure CLI session (including Secure Copy) **4.1.2** SSL for secure HTTP session
Switch management	Local authentication database Remote authentication RADIUS, TACACS+, LDAP & ACE servers **4.2.19**
802.1X/Device Authentication	Supported platform: OS6800, OS6850, and OS9000 There are 4 levels of 802.1x/device classification: -Basic 802.1x port. Only successful authenticated 802.1x devices are allowed in the network -Basic 802.1x port + fail authentication policies. Only 802.1x capable devices are allowed in the network. These policies allow the failed authenticated 802.1x devices to access non-secured (or non authenticated) VLANs -802.1x + non supplicant policies without Mac authentication. Non 802.1x devices are allowed on non-secured VLANs according to the non-supplicant policies. -802.1x + non supplicant policies with Mac authentication. In this mode, the non 802.1x devices will follow either the “non-supplicant authentication pass policies” when the Mac authentication is successful or the “non-supplicant authentication fail policies” when the Mac authentication failed The open-unique and open-global options are no longer applicable. Device Authentication: Maximum number of supplicants / non-supplicant users per system: 1024

	Maximum number of non-supplicant users per port: 1024 Maximum number of supplicant users per port: 253 Maximum combined number of supplicant and non-supplicant users per port: 1024 The system supports up to 1024 authenticated/mobile Mac-addresses. Supported/non-supported mobile rule on device authentication: 1. Support rule per tagged/untagged packet type. Mac rule – apply on UNTAGGED packet IP subnet rule – apply on UNTAGGED packet Protocol rule – apply on UNTAGGED packet Port-protocol binding rule– apply on UNTAGGED packet Mac-port binding rule – apply on UNTAGGED packet Mac-IP-port binding rule– apply on UNTAGGED packet Mobile-tag – apply on TAGGED packet * Mobile tag only apply on tagged packets, all other rules apply on untagged packet. 2. DHCP related mobile rules are not supported with device authentication (i.e. supplicant/non-supplicant cases) DHCP generic rule DHCP port rule DHCP Mac / Mac range rule Device authentication with Alcatel.Lucent IP phone: Alcatel.Lucent Dynamic IP phone has 3 modes: 1.Untagged dynamic Packet is always untagged. 2.Tagged dynamic Packet is always tagged based on administrator config on phone. 3.Alcatel.Lucent dynamic First packet is untagged, second packet onward is tagged.
ACLMAN **4.2.16**	ACLMAN is a function of the QoS subsystem in AOS. ACLMAN allows a network administrator to manage ACLs using default industry standard syntax on Alcatel.Lucent switches. To enforce the ACLs, ACLMAN translates default industry standard syntax into Alcatel.Lucent QoS filtering policies in a manner transparent to the ACLMAN user. ACLMAN provides the following: · The ability to import text files from flash containing default industry standard ACL syntax · An interactive shell emulating the default industry standard CLI ACL command syntax ACLMAN supports the following default industry standard ACL types: · Standard ACLs · Extended ACLs · Numbered ACLs · Named ACLs These are the limitations for the 6.1.2.R03 release. - Only supported on the OS6850 Series - No stacking support - ACLMAN is restricted by the same number of rule limitations that QoS supports - ACL names are limited to 16 characters
Management	
Configuration Mode	Command Line Interface (CLI) Telnet/SSHv2 for remote CLI access Web-based (HTTP/HTTPS) SNMPv1/v2c/v3 for complete NMS integration **4.1.3**
Management Access types	Serial Console port for local & remote (modem dial up) access (RJ45) Out-of-band Ethernet access (10/100/1000RJ45) In-band Ethernet access
System Maintenance	Port Mirroring (one-to-one, many-to-one) RMON (Remote Monitoring): Statistics, History, Alarm & Events **4.1.3** sFlow Local & Remote logging (Syslog) **4.2.17** Detailed Statistics / Alarm / Debug information per process L3 OAM (ICMP Ping and Traceroute) NTP (Network Time Protocol) Internal flash (Compact Flash) to feature: • Working Directory • Certified Directory
System file Transfer	XModem FTP (File Transfer Protocol) **4.1.4**
Working and Certified Directories	OmniSwitch 9000 Series switches are shipped with 128 MB of flash memory. This memory is used to store files, including boot and image files that are used for switch operations. The /flash directory contains two subdirectories: /working and /certified. These directories work

	together to provide the image rollback resiliency feature. Image rollback allows the switch to return to a prior "last known good" version of software in the event of a system software problem. The /flash/working directory is intended for software that is still being configured for your network. Changes made while configuring your switch are saved to the boot.cfg file in the /flash/working directory. Once the /flash/working directory's configuration and image files are road-tested and considered valid and reliable for your network, they can be copied to the /flash/certified directory. The software in the /flash/certified directory should be treated as the "gold master" for the switch. When you place configuration and image files in this directory, you are "certifying" them as tested and reliable. If the switch is running from the /flash/working directory and experiences a software problem, it will "roll back" to the last known good software in the /flash/certified directory on the next reboot.
Using the WebView	OmniSwitch 9000 switches can be configured and monitored using WebView, Alcatel.Lucent's Web-based device management tool. WebView software is pre-installed in the switch; you are not required to load additional software. Note. Although WebView software is pre-installed, you must first enable HTTP sessions for your switch before you can log in. WebView has been tested on the following Web browsers: • Internet Explorer 6.0 for Windows 2000, Windows NT, and Windows XP • Netscape 4.79 for Solaris 2.8, and HP-UX 11.0 • Netscape 7.1 for Windows 2000, Windows NT, and Solaris 2.8
Port Disable	You can configure a "Port Disable" rule to administratively disable an interface when matching a policy rule. To make the interface operational again, the port must be unplugged/plugged back or disabled/enabled using "interfaces s/p admin down" and "interfaces s/p admin up". Also, a SNMP trap will be sent when an interface goes down when matching a port disable rule.
SNMP Traps	A "pktDrop' SNMP trap will be sent out to the SNMP station when a port goes down because of a user-port shutdown profile or a port disable rule.
Ethernet Services OAM	Ethernet services can be offered over multiple types of transport using a variety of tunneling technologies. In all such layered models, it is important to provide basic OAM capabilities in each layer of the hierarchy. Ethernet Services OAM addresses the OAM functionality in the Ethernet Service (ETH) layer, which remains independent of the underlying TRAN layer(s), each of which may have its own OAM capability. The requirements of OAM functions for the ETH layer focus on monitored parameters e.g. connectivity, delay, delay variation (jitter) and status monitoring. The Ethernet service interface is considered to be the OAM source and termination of ETH layer OAM. In particular, the Ethernet service interface on each device is assumed to have a MAC address that can be used for OAM packet addressing. **Feature to be supported with AOS 6.3.1R01**
SFLOW **4.2.17**	SFlow is a sampling technology embedded within switches/routers defined in RFC 3176. It provides the ability to monitor the traffic flows. It requires an sFlow Agent running in the Switch/Router and a sFlow collector which receives and analyses the monitored data. SFlow agent running on the OS6850, combines interface counters and traffic flow (packet) samples on all the configured interfaces into sFlow Datagrams that are sent across the network to an sFlow collector (3rd Party software). Packet sampling is done in hardware and is non-CPU intensive. Current release (6.1.3r01) will not support IPv6 as Collector. **The switch sends the first 128 bytes of the sampled packet from which the entire layer 2/3/4 information can be extracted by the receiver.** This could include: - Source/Destination Mac address - Source IP/ Destination IP - Source/Destination TCP/UDP/ICMP port - Source/Destination Physical port (Gigabit Port) - IPv4/IPv6 - RIP/OSPF/BGP/PIM-SM/DM (OK, but if this information falls within the first 128 Bytes of the packet) - VLAN - QoS 802.1Q, ToS and DiffServ (DSCP) - Data Payload (OK, but if this information falls within the first 128 Bytes of the packet) - Others (If this information falls within the first 128 Bytes of the packet) Given an IP Address the SFLOW sampling information can be sent to a Collector such as the InMon and/or the Crannog.
SFLOW Back-off Algorithm	Since the CPU of switch is involved in the datagram processing, there is a built in back-off algorithm which will automatically adjust the sampling rate in the case of CPU congestion on switch. This back-off mechanism is not user-configurable in Release 6.1.3r01. If CPU is congested it automatically continues to double the sampling rate, and will continue to do so up to a very low rate of 1 sample in 2147483647 (2exp31)-1. For a 1Gig interface, the bit rate is 1,000,000,000 bits per second. The back-off algorithm is designed to take effect when the sample rate exceeds 10 samples per second on any interface. Since each sample is configured by default for 128bytes this is 10x128x8 or 10 samples/sec x 1024 bits/sample or 10x1024 bps 1Gbps / 10x1024 bps = 97656 sampling rate. Sampling with all available slot/ports at 10G wire-rates on OS9000 and all ports at 1G on the OS6850

<table>
<tr><td></td><td>keep backing-off up to 2,147,483,647 and stay fixed at this value until the traffic generation is halted or reduced. That is even running only one 1G interface at wire rate on the OS6850 will back-off to 2147483647 and stay at this (maximum, safe) sampling rate.
Recommended sampling rates for various speeds at various load:
<table><tr><th></th><th colspan="3">Sampling Rates</th></tr><tr><th>Link Speed</th><th>Light Load</th><th>Medium Load</th><th>Heavy Load</th></tr><tr><td>10Mb/s</td><td>256</td><td>512</td><td>8192*</td></tr><tr><td>100Mb/s</td><td>512</td><td>1024</td><td>65536*</td></tr><tr><td>1Gb/s</td><td>1024</td><td>2048</td><td>Max*</td></tr><tr><td>10Gb/s</td><td>2048</td><td>4096</td><td>Max*</td></tr></table>*8192 is the empirical value found in the lab for 10Mbs, 65536 for 100 Mbps
*Max: because the OS6850 always backs-off to a max sampling rate of 2147483647 for wire rate at these rates. All other values are those recommended by Inmon. Whatever the configured sampling rate, the back-off mechanism will set the ‘meanskipcount’ higher or lower depending on what is the ‘unaffecting’ sampling rate for the CPU.</td></tr>
<tr><td>TACACS+</td><td>Supported platform: OS6800, OS6850, and OS9000
Release 6.1.3.R01 is the first release to support TACACS+ AAA.
AOS implementation is based on the Tacacs+ Protocol: draft-grant-tacacs-02.txt, January 1997.
Overview:
ASA or Authenticated Switch Access to AOS OmniSwitch running 6.1.3.R01 can be configured to add servers and forward AAA requests to TACACS+. TACACS+ servers are configured similar to RADIUS or LDAP servers; however, (MD5) encryption key is optional.
AAA authentication and accounting services must be configured to point to the desired TACACS+ server. It is possible to set authentication and authorization to one TACACS+ server and accounting requests to a different server.
The number of configurable servers and fail over to second server is uniform across all AAA server types: Up to 4 servers can be configured and all queries will be sent to the 1st server only. If 1st server is online and user exists on 2nd server, the result will be failed authentication. If the 1st server is down, authentication and authorization requests will only be sent to “next available” server. If all servers are down, all logins will fail.
Different AAA services can be configured to query different authentication servers. All services may use a common authentication protocol or mix of supported protocols: Telnet service may be configured to query RADIUS while http/ftp may be configured to query TACACS+. Or all may query RADIUS. Or all may query TACACS+. In all cases accounting server protocol must match authentication/authorization server protocol.
AOS TACACS+ does not support authentication for network or windows domain access. Only AOS switch access with Partition Management type domain family attribute/value pairs is supported. This to say different users or groups of users may be assigned various levels of AOS switch management privileges.
The TACACS+ servers run as an external server on Unix or Windows. We have tested with CISCO TACACS+ freeware for Unix and Cisco’s Secure ACSv4.0
TACACS+ uses TCP instead of UDP. Each login and supported command is queried back to the server for authorization.
TACACS+ configuration is fully supported with AOS WebView.
Notes:
•Tacacs+ supports Authenticated Switch Access and cannot be used for user authentication.
•Authentication and Authorization operations are combined together and cannot be performed independently. This implies that when Tacacs+ authentication is enabled, Tacacs+ authorization is also enabled. Disabling Tacacs+ authentication automatically disables authorization.
•A maximum of 50 simultaneous Tacacs+ sessions can be supported, when no other authentication mechanism is activated. This is a limit enforced by the AAA application.</td></tr>
<tr><td>Power over Ethernet</td><td>The Standard in brief

• In IEEE 802.3af standard, POE transmits power over the same pair as the data. This method is called the resistive detection method.
• In non-802.3af or pre-802.3af standard, POE transmits power over a spare pair (not the same pair as the data). This method is called the capacitor detection method.

Max power per port

• The max power per port is 18 watts for OS9000. Using 350 milliamps in the standard to calculate max power, this is based on tight tolerances (+-0.5) for OS9000 POE power supplies (Vmain) at 52 volts.</td></tr>
</table>

	Max power per blade • OS9000 lanpower is load-shared among all of the GNI-P24 NIs in the chassis, each NI having 210W max per blade for lanpower. Up to 4 power supplies (525W x 4) of 2100W max in the power shelf is available for the entire chassis (please note that up to 2400watts of PoE power will be supported in a future release). Depending on how many GNI-P24s in the chassis and how much power all the NIs required, power supply redundancy can be defined as having at least one power supply more than the power requirement. Note that a DB25 female-male power cable is needed in order to connect between the chassis and the power supply. Each power cable must be plugged in to the corresponding connector ID between both the power shelf and switch chassis because the i2c reading retrieves power supply information accordingly. • OS9600 can also support either a 510W or 360W (normally used with the OS6850) power supply and it is load-shared among all of the GNI-P24 NIs in the chassis. Each NI can only have 210W max power blade for lanpower. For OS9600 chassis, there is no redundancy support using "OS6850" power supply because only one power supply can be inserted into the chassis. Note that a DB25 male-male power cable is needed in order to connect between the chassis and the power supply. **Lanpower Priority** • For port-priority, both the OS9000 is set to "low" by default in all the ports. Therefore the lowered-numbered ports always have a higher-precedence of retaining lanpower when there is insufficient power for all the ports. In order for higher-numbered ports to have a higher priority, use the CLI command to set the port priority higher "lanpower <slot/port> priority <low/high/critical>". • For slot-priority, OS9000 is set to low by default in all the slots. The lower-numbered slot has a higher precedence of retaining lanpower when there is insufficient power for all the slots. As a result, slot-priority will override port-priority setting no matter what. In order for higher-numbered slot to have a higher priority, use the CLI command to set the slot priority higher "lanpower <slot> priority <low/high/critical>".
Software	
Capability Maturity Model (CMM)	Alcatel.Lucent's Software Engineering Institute (SEI) Capability Maturity Model (CMM) rating for software processes meets the Level-2 (CMM-level-2) requirements.
The Ethernet software	The Ethernet software is responsible for a variety of functions that support the Ethernet, Gigabit Ethernet and 10Gigabit Ethernet ports on OmniSwitch 9000 Series switches. These functions include diagnostics, software loading, initialization, and configuration of line parameters, gathering statistics, and responding to administrative requests from SNMP or CLI.

Operating Systems

Wind River's VxWorks multi-tasking O/S version 5.4 with a Kernel version 2.5. Alcatel.Lucent O/S – AOS (Alcatel.Lucent's Operating Systems).
The AOS is uploaded onto the Flash memory. The advantage of this switch running the AOS is that it is managed using the same interface as with the rest of the Alcatel.Lucent AOS switching & routing platforms. The AOS on the OS9000 platforms provides support for the majority of the features of the larger modular platforms including layer-3 unicast routing using RIPv1&v2, VRRP, or OSPFv2. Group mobility and authenticated VLANs as well as QoS and ACL functionality are supported making the OS9000 a highly functional solution for the core of the network.

Software

Each OmniSwitch 9000 Chassis is shipped with base software. All advanced features are also included in the base software.

Authenticated Services Software	
OS9000-SW-SBR-N	OS9000 Auth. SBR-MS SW w/MD5, RC4, MD4, DES. OmniSwitch 9000 Authentication Services software bundled with Funk Software's Steel-Belted Radius Enterprise Edition for Microsoft Windows. An annual maintenance contract, 801159-00 (SER-SBR), for Funk SBR must be purchased with this reference.
OS9000-SW-SBR-S	OS9000 Auth. SBR-Sun SW w/MD5, RC4, MD4, DES. OmniSwitch 9000 Authentication Services software bundled with Funk Software's Steel-Belted Radius Enterprise Edition for Sun Solaris. An annual maintenance contract, 801159-00 (SER-SBR), for Funk SBR must be purchased with this reference.

Simplified Manageability

ecognizing a great demand in the marketplace and customers expectation for a level of synergism in the network management, Alcatel.Lucent has developed a comprehensive, unified, and simplified network and switch management solutions for its full array of networking products including the AOS OmniSwitch product family. The OS9000 switch and network management industry proven features offers the network administrators ease-of-use and ease of management. The following is only a highlight of the advanced network and switch management features supported by the OmniSwitch 9000 Series:

- Diagnosing Switch problems:
  - Port Mirroring
  - RMON: Supports RFC 2819 RMON group (1-Statistics, 2-History, 3-Alarm, and 9-Events)
  - Switch Health
  - Monitoring Memory Tools
  - Switch Logging
- Authentication or AAA Servers
- Policy Servers
- Dual image and dual configuration file storage provides backup **4.2.18**
- Intuitive Alcatel.Lucent CLI for familiar interface and reduced training costs
- Easy to use point and click web based element manager with built-in help for easy configuration of new technology features
- Remote telnet management or secure shell
- Port based, port mirroring for troubleshooting, supports four sessions with four source to one destination configuration.
- Human readable ASCII based config files for offline editing and bulk configuration
- IGMPv1/v2/v3 snooping to optimize multicast traffic
- BootP/DHCP client allows auto-config of switch IP information to simplify deployment
- Auto-negotiating 10/100/1000 ports automatically configure port speed and duplex setting
- Auto MDI/MDIX automatically configures transmit and receive signals to support straight thru and crossover cabling
- DHCP relay to forward client requests to a DHCP server **4.1.5**
- SNMPv1/v2/v3
- Integration with SNMP manager OmniVista for network wide management
- System event log
- Network Time Protocol (NTP) for network wide time synchronization
- Alcatel.Lucent Interswitch Protocols (AIP) **4.1.6**
  - AMAP: Alcatel.Lucent Mapping Adjacency Protocol (AMAP) for building topology maps within OmniVista
  - GMAP

Alcatel.Lucent Interswitch Protocols (AIP)	Alcatel.Lucent Interswitch Protocols (AIP) is used to discover adjacent switches and retain mobile port information across switches. The following protocols are supported: • Alcatel.Lucent Mapping Adjacency Protocol (AMAP), which is used to discover the topology of OmniSwitches and OmniSwitch/Routers (Omni S/R). • Group Mobility Advertisement Protocol (GMAP), which is used to retain learned mobile port and protocol information. These protocols are independent of each other and perform separate functions.
Interswitch Protocol (AMAP)	Alcatel.Lucent Interswitch Protocols (AIP) is used to discover adjacent switches and retain mobile port information across switches. By default, AMAP is not enabled. The Alcatel.Lucent Mapping Adjacency Protocol (AMAP) is used to discover the network topology of OmniSwitch, switches in a particular installation. Using this protocol, each switch determines which OmniSwitch; Omni S/R and/or OmniAccess switches are adjacent to it by sending and responding to Hello update packets. For the purposes of AMAP, adjacent switches are those that: • Have a Spanning Tree path between them • Do not have any switch between them on the Spanning Tree path that has AMAP enabled
Authentication Servers or AAA servers (authentication, authorization, and accounting)	Authentication servers are sometimes referred to as AAA servers (authentication, authorization, and accounting). These servers are used for storing information about users who want to manage the switch (Authenticated Switch Access) and users who need access to a particular VLAN(s) (Authenticated VLANs). RADIUS or LDAP servers may be used for Authenticated Switch Access and/or Authenticated VLANs. Another type of server, SecurID's ACE/Server, may be used for authenticated switch access only; the ACE/Server is an authentication-only server (no authorization or accounting). Only RADIUS servers are supported for 802.1X Port-Based Network Access Control. **Authentication Servers Specifications:** RADIUS RFCs Supported: ▪ RFC 2865–Remote Authentication Dial In User Service (RADIUS) ▪ RFC 2866–RADIUS Accounting ▪ RFC 2867–RADIUS Accounting Modifications for Tunnel Protocol Support ▪ RFC 2868–RADIUS Attributes for Tunnel Protocol Support ▪ RFC 2809–Implementation of L2TP Compulsory Tunneling via RADIUS ▪ RFC 2869–RADIUS Extensions ▪ RFC 2548–Microsoft Vendor-specific RADIUS Attributes ▪ RFC 2882–Network Access Servers Requirements: Extended RADIUS Practices LDAP RFCs Supported: ▪ RFC 1789–Connectionless Lightweight X.5000 Directory Access Protocol ▪ RFC 2247–Using Domains in LDAP/X.500 Distinguished Names

	▪ RFC 2251–Lightweight Directory Access Protocol (v3) ▪ RFC 2252–Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions ▪ RFC 2253–Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names ▪ RFC 2254–The String Representation of LDAP Search Filters ▪ RFC 2256–A Summary of the X.500 (96) User Schema for Use with LDAPv3 Other RFCs: ▪ RFC 2574–User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) ▪ RFC 2924–Accounting Attributes and Record Formats ▪ RFC 2975–Introduction to Accounting Management ▪ RFC 2989–Criteria for Evaluating AAA Protocols for Network Access
Authentication Servers	Maximum number of authentication servers in single authority mode: ▪ 4 (not including any backup servers) Maximum number of authentication servers in multiple authority mode: ▪ 4 per VLAN (not including any backup servers) Maximum number of servers per Authenticated Switch Access type: ▪ 4 (not including any backup servers) CLI Command Prefix Recognition: ▪ The aaa radius-server and aaa ldap-server commands support prefix recognition.
ACE/Server	An external ACE/Server may be used for authenticated switch access. It cannot be used for Layer 2 authentication or for policy management. Attributes are not supported on ACE/Servers. These values must be configured on the switch through the user commands. Since an ACE/Server does not store or send user privilege information to the switch, the switch determines user privileges for Secur/ID logins. When a user attempts to log into the switch, the user ID and password is sent to the ACE/Server. The server determines whether the login is valid. If the login is valid, the user privileges must be determined. The switch checks its user database for the user's privileges. If the user is not in the database, the switch uses the default privilege, which is determined by the default user account. There are no server-specific parameters that must be configured for the switch to communicate with an attached ACE/Server; however, you must FTP the sdconf.rec file from the server to the switch's/network directory. This file is required so that the switch will know the IP address of the ACE/Server. The ACE client in the switch is version 4.1; it does not support the replicating and locking feature of ACE 5.0, but it may be used with an ACE 5.0 server if a legacy configuration file is loaded on the server. The legacy configuration must specify authentication to two specific servers (master and slave). The ACE/Server generates "secrets" that it sends to clients for authentication. While you cannot configure the secret on the switch, you can clear it. The secret may need to be cleared because the server and the switch get out of synch.
RADIUS Servers	RADIUS is a standard authentication and accounting protocol defined in RFC 2865 and RFC 2866. A built-in RADIUS client is available in the switch. A RADIUS server that supports Vendor Specific Attributes (VSAs) is required. The Alcatel.Lucent attributes may include VLAN information, time-of-day, or slot/port restrictions. RADIUS Server Attributes: RADIUS servers and RADIUS accounting servers are configured with particular attributes defined in RFC 2138 and RFC 2139, respectively. These attributes carry specific authentication, authorization, and configuration details about RADIUS requests to and replies from the server. For a complete list of attributes (standard, and vendor-specific) and how to configure them on the server, please refer to the Users Manual.
Lightweight Directory Access Protocol (LDAP)	Lightweight Directory Access Protocol (LDAP) is a standard directory server protocol. The LDAP client in the switch is based on several RFCs: 1798, 2247, 2251, 2252, 2253, 2254, 2255, and 2256. The protocol was developed as a way to use directory services over TCP/IP and to simplify the directory access protocol (DAP) defined as part of the Open Systems Interconnection (OSI) effort. Originally it was a front-end for X.500 DAP. The protocol synchronizes and governs the communications between the LDAP client and the LDAP server. The protocol also dictates how its databases of information, which are normally stored in hierarchical form, are searched, from the root directory down to distinct entries. In addition, LDAP has its own format that permits LDAP-enabled Web browsers to perform directory searches over TCP/IP. For a complete list of attributes (vendor-specific) and how to configure them on the server, please refer to the Users Manual.
Policy Servers (Policy Server Management)	Quality of Service (QoS) policies that are configured through Alcatel.Lucent's PolicyView network management application are stored on a Lightweight Directory Access Protocol (LDAP) server. PolicyView is an OmniVista application that runs on an attached workstation. **Policy Server Specifications:** LDAP Policy Servers RFCs Supported: ▪ RFC 2251–Lightweight Directory Access Protocol (v3) ▪ RFC 3060–Policy Core Information Model—Version 1 Specification Maximum number of policy servers (supported on the switch): 4 Maximum number of policy servers (supported by PolicyView): 1 Policy servers use the Lightweight Directory Access Protocol (LDAP) to store policies that are configured through Alcatel.Lucent's PolicyView network management application. PolicyView is an OmniVista application that runs on an attached workstation.

	The Lightweight Directory Access Protocol (LDAP) is a standard directory server protocol. The LDAP policy server client in the switch is based on RFC 2251. Currently, only LDAP servers are supported for policy management. The switch communicates with the LDAP server to download and manage LDAP policies. When the policy server is connected to the switch, the switch is automatically configured to communicate with the server to download and manage policies created by the PolicyView application. There is no required user configuration. (Note that the LDAP policy server is automatically installed when the Policy-View application is installed.) Note. The switch has separate mechanisms for managing QoS policies stored on an LDAP server and QoS policies configured directly on the switch.